

GROUPWARE
FOR
THE
MACINTOSH[®]

A COMPLETE
GUIDE TO
COLLABORATIVE
COMPUTING

Michael Fraase

COMPUTERS/SOFTWARE (SPRING 1991)

Other BUSINESS ONE IRWIN Titles of Interest to You

DESKTOP PUBLISHING SUCCESS

How to Start and Run a Desktop Publishing Business

Felix Kramer and Maggie Lovaas

Common-sense solutions and strategies for succeeding in a DTP career. A valuable guide for entrepreneurs and for those establishing and maintaining a DTP department.

ISBN: 1-55623-424-4 \$27.50 (paper)

SPECTACULAR COMPUTER CRIMES

What They Are and How They Cost American Business Half a Billion Dollars a Year!

Buck BloomBecker

Mandatory reading for all professionals who rely on computers to handle valuable data. Recounting actual cases, BloomBecker illustrates ways that computer users are vulnerable to computer crime and how they can protect themselves against it.

ISBN: 1-55623-256-X \$24.95

FARALLON'S MEDIATRACKS™

The Ultimate Training Tool

Michael Fraase

MediaTracks is ideal for training because users can create an entire interactive personnel training session in an hour or less. Fraase guides you through this powerful new management tool in an easy to read and easy to remember narrative style.

ISBN: 1-55623-426-0 \$37.50 (paper)

NEW LIFE FOR OLD PCs

How to Keep Your Company's Computers from Becoming Obsolete

Alfred Poor

Shows you how to extend your computer investment dollars. Poor describes strategies for effectively salvaging, converting, and upgrading your PC. Includes a user-friendly software disk with a practical checklist.

ISBN: 1-55623-427-9 \$29.95 (paper)

WEALTHBUILDER™ BY MONEY MAGAZINE

Taking Control of Your Financial Destiny

Donald R. Woodwell and Doris Woodwell

WealthBuilder by Money Magazine is an easy to use financial planning software program. The Woodwells show you how to use it to plan all aspects of your financial goals including tax savings and home purchasing.

ISBN: 1-55623-441-4 \$24.95 (paper)

POWER TO GO

Bruce Brown's Guide to Getting the Most from Portable Computers

Bruce Brown

Provides portable computer owners with essential information for determining the right software/hardware mix. Extensive vendor lists help you quickly compare and choose.

ISBN: 1-55623-439-2 \$24.95 (paper)

852

Groupware for the Macintosh

***A Complete Guide to
Collaborative Computing***

Groupware for the Macintosh

***A Complete Guide to
Collaborative Computing***

Michael Fraase

**BUSINESS ONE IRWIN
Homewood, Illinois 60430**



The information in this book is distributed on an "As Is" basis, without warranty. Neither the author nor Business One Irwin shall have any liability to customer or any other person or entity with respect to any liability, loss, or damage caused or alleged to be caused directly or indirectly by the programs contained herein. This includes, but is not limited to, interruption of service, loss of data, loss of business or anticipatory profits, or consequential damages from the use of the programs.

Business One Irwin professional books are available for bulk sales at quantity discounts. For information, please contact Marketing Manager, Professional Books Group, Business One Irwin, 1818 Ridge Road, Homewood, IL 60430.

© RICHARD D. IRWIN, INC., 1991

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold with the understanding that neither the author nor the publisher is engaged in rendering legal, accounting, or other professional service. If legal advice or other expert assistance is required, the services of a competent professional person should be sought.

From a Declaration of Principles jointly adopted by a Committee of the American Bar Association and a Committee of Publishers.

Sponsoring editor: Susan Stevens, Ph.D.
Project editor: Karen J. Nelson
Production manager: Irene H. Sotiroff
Printer: The Book Press, Inc.

Library of Congress Cataloging-in-Publication Data

Fraase, Michael.

Groupware for the Macintosh : a complete guide to collaborative computing / by Michael Fraase.

p. cm.

Includes index.

ISBN 1-55623-521-6

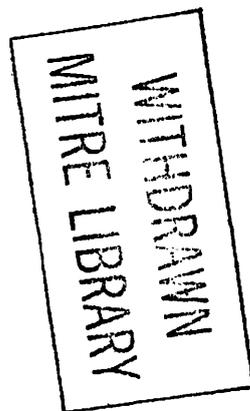
1. Work groups—Data processing.
2. Business—Data processing.
3. Business—Communication systems.
4. Information technology.

I. Title.

HD66.F73 1992

658.4'036'028546—dc20

91-9949



Printed in the United States of America

1 2 3 4 5 6 7 8 9 0 BP 8 7 6 5 4 3 2 1

Trademark Credits

All terms mentioned in this work that are known to be trademarks or service marks are listed below. In addition, terms suspected of being trademarks or service marks have been appropriately capitalized. Use of a term in this work should not be regarded as affecting the validity of any trademark or service mark.

Apple, the Apple logo, AppleTalk, HyperCard, ImageWriter, LaserWriter, Macintosh, and MacTerminal are registered trademarks of Apple Computer Inc.

AppleCD, Apple Desktop Bus, AppleFax, AppleMouse, AppleShare, AppleShare Print Server, A/UX, EtherTalk, Finder, HyperTalk, LaserShare, LocalTalk, and MultiFinder are trademarks of Apple Computer Inc.

Arts & Farces, the Arts & Farces logo, and Arts & Farces Review are trademarks of Arts & Farces.

CheckNet, Liaison, NetAtlas, PhoneNET, Timbuktu, Timbuktu/Remote, and Traffic Watch II are trademarks of Farallon Computing Inc.

CompuServe is a registered trademark of CompuServe Inc.

DataClub is a trademark of International Business Software.

Ethernet is a registered trademark of Xerox Corporation.

GatorBox and GatorShare are trademarks of Cayman Systems Inc.

GoFER is a trademark of Microlytics Inc.

GraceLAN is a trademark of Technology Works Inc.

Inspiration is a trademark of Ceres Software Inc.

MacLink Plus is a trademark of DataViz Inc.

MacProject is a registered trademark of Claris Corporation.

MarkUp is a trademark of Mainstay Software.

Microsoft and MS-DOS are registered trademarks of Microsoft Corporation.

Microsoft Mail is a trademark of Microsoft Corporation.

MORE is a trademark of Symantec Corporation.

NetUpdater is a trademark of MDG Computer Services.

NuBus is a trademark of Texas Instruments.

On Location is a trademark of On Technology Inc.

PC-DOS is a registered trademark of International Business Machines Corporation.

Personal Server Network is a trademark of Information Presentation Technologies Inc.

PostScript is a registered trademark of Adobe Systems Inc.

QuickMail is a trademark of CE Software Inc.

Retrospect is a trademark of Dantz Development Corp.

TeleFinder is a trademark of Spider Island Software.

Timeslips is a trademark of Timeslips Corporation.

UNIX is a registered trademark of AT&T Information Systems.

Preface

Groupware, collaborative computing, and computer-supported cooperative work (CSCW) are terms you'll be hearing more and more in the next few years. They all mean the same thing: using computers to work in a collaborative workgroup setting. Groupware can be defined as a group of technologies, techniques, and services designed to help people collaborate more effectively, productively, and creatively. Groupware can consist of hardware, software, services, and support.

Computers have returned philosophy to an integrated part of everyday life. When we work with computers we are forced to ask questions about the world around us:

- What is thought?
- What is intelligence?
- What is the mind?
- What is reality?

Groupware tools cause us to ask these questions from several different viewpoints because they enable us to work collaboratively, with others, beyond the normal constraints of time and space that we have been taught to accept. These tools allow us to collaborate with someone who is thousands of miles away, just as if he or she were in the next office down the hall.

These tools let us try new things without the fear of failure. The reason desktop publishing has become so successful is

because the computer gives us permission to fail. We try a multitude of designs without having to “marry” ourselves to any single idea until we reach one that we recognize is “right.” Then we can refine that design as never before.

Our abilities grow because the computer frees us to take new risks without committing to the outcome. Our abilities grow because the computer extends not only what we can *do*, but also what we can *imagine*.

Groupware tools allow us to extend that willingness to fail to the workgroup. Collaborations can be more fruitful with the use of these tools, and the group process itself is redefined.

The tools and techniques described in this book will help you collaborate more effectively by helping you manage information—both shared and private—more efficiently.

Who Can Benefit From This Book

If you use a Macintosh in a workgroup setting, or if you collaborate with at least one other person you can benefit from this book.

The most important development of our lifetime has been the replacement of the second-wave, industrial economy with one based on information and knowledge. Instead of creating wealth with muscle, as we did in the industrial economy, we now create wealth by manipulating symbols—by converting information into knowledge. More and more we’re doing this in workgroups rather than individually.

We all collaborate. Maybe you aren’t part of a permanent workgroup (and as the economy continues to shift, the notion of a “permanent” anything in the workplace will continue to fade), but you collaborate with others nonetheless.

I can think of no profession that is perceived as being more solitary than that of the professional writer. This is a misperception, however, and this book is a good example. I wrote it, and my name is on the cover, but it was anything but a solitary effort. I spoke with literally hundreds of people during my research and writing cycles for this book. I worked with an editor, a copyeditor, a promotion department, and a service bureau to produce the camera-ready pages that you hold in your hands. It was a patchwork of various collaborations, some of which endured throughout the project and some of which were quite brief. In any case, this book would not exist if not for those collaborations.

I suspect that your work environment has at least that much in common with mine. If so, you can benefit from this book.

Navigating This Book

The chapters in this book were designed, as much as possible, to be free-standing. *Groupware for the Macintosh: A Complete Guide to Collaborative Computing* assumes that you have an underlying knowledge of basic Macintosh techniques: launching applications, selecting objects, creating folders, saving files, and the like. If you don't, refer to Apple's excellent guides that came with your computer.

Chapter One, Introduction and Overview, provides a brief overview of the underlying concepts of collaborative computing principles, technologies, and techniques.

Chapter Two, Driving Forces of Groupware, offers background material on our shifting economy and why groupware is an important issue for businesses and individuals alike. An in-depth discussion of both classic and modern information theory is provided, and the basic premises of workgroup collaboration and the collaborative process are explored.

Chapter Three, Collaborative Computing With Groupware, details networking fundamentals and evaluations of currently available network media. The chapter also provides information on planning, managing, and maintaining a networked work environment, including collaborative computing in the mixed platform environment.

Chapter Four, Groupware Implementations, offers an overview to the different approaches to groupware identified by Robert Johansen. The chapter also explores existing large-scale collaborative environments with a focus on Xerox's Colab. Complete evaluations of the leading groupware products for the Macintosh are also included in this chapter.

Chapter Five, Tutorials, provides a series of step-by-step tutorials for establishing and managing a collaborative workgroup environment.

Chapter Six, Creating the Collaborative Environment, explores the human side of the collaborative process. This chapter offers suggestions for enhancing the effectiveness of collaborative activities within the workgroup setting.

Chapter Seven, Collaborative Hypertext and Hypermedia, examines the role of the new interactive media in the collaborative environment.

Chapter Eight, Problems With Groupware, examines the problems associated with collaborative computing and the workgroup environment.

Chapter Nine, Future Groupware, provides a brief glimpse at the future of groupware and its logical extensions.

The Appendices include a glossary, product source listing, and bibliography.

Michael Fraase
Saint Paul, Minnesota
March 1991

Table of Contents

Chapter One: Introduction and Overview

Vannevar Bush's Memex	19
Douglas Engelbart's Augmentation Work.....	23
NLS and Augment.....	24
The Knowledge Workshop	25
Augmentation	27
A Conceptual Framework and Augmentation Means.....	28
Ted Nelson's Collaborative Computing Work.....	30
Xanadu.....	30
Thinkertoys and Super Virtualities	34

Chapter Two: Driving Forces of Groupware

The Information Economy	37
A New Hierarchy	39
<i>Information Expands as it is Used</i>	40
<i>Information is Compressible</i>	40
<i>Information is Transportable</i>	41
<i>Information Wants to Leak</i>	41
<i>Information Wants to be Shared</i>	41
The Collapse of the Information Float	42
Stanley Davis's Four Earmarks of the Information Economy.....	43
Information and the "Time is Money" Dichotomy.....	43
Claude Shannon's Information Theory	45
<i>Redundancy</i>	45
Groupware Technologies and the Mass-Customization of Information.....	47
Information and the Anxiety it Causes	48
<i>Richard Saul Wurman's Information Immediacy</i>	49

Workgroup Collaboration.....	52
Interpersonal Computing	55
Why Collaborate?.....	57
The Collaborative Process.....	59
Implications for Interpersonal Interaction.....	61

Chapter Three: Collaborative Computing With Groupware

Networking Basics.....	63
The Network as Electronic Highway	64
Hardware Components	64
Software Components	68
Network Architecture	68
Communications Protocols	68
The ISO/OSI Model	69
Topology	70
Bus Topology.....	70
Ring Topology.....	71
Star Topology.....	72
Categories of Networks	73
Internetworks	75
Nodes and Addressing.....	79
Packets.....	80
Access Methods	80
Transmission Methods	81
Plugging In	82
Built-in Circuitry.....	82
Add-on Circuit Board	82
Modem	82
Error Correction	83
Data Compression.....	83
Microcom Networking Protocol (MNP).....	84
Modulation Standards.....	85
Network Services	86
File Service.....	86
Dedicated file server	87
Nondedicated file server	88
Distributed file server.....	90

Electronic Mail	91
Printer Service	92
Network Standards	93
The Future of Networking	93
ISDN	94
<i>Basic Rate Interface (BRI)</i>	94
<i>Primary Rate Interface (PRI)</i>	95
Potential ISDN Integrated Services	97
<i>Screen-Based Telephony</i>	97
<i>Imaging</i>	98
<i>Resource Sharing</i>	98
<i>Things We Can't Even Imagine Yet</i>	99
Fiber Optic	100
Managing the Network	101
Assessing your Networking Needs	101
<i>Inventory Your Current Needs</i>	102
<i>Determine Network Service Needs</i>	102
<i>Evaluate Solutions</i>	103
<i>Anticipate and Prepare for Growth</i>	103
<i>Decide on the Physical Layout</i>	104
Administering the Network	104
Collaborative Computing in the Mixed Platform Network	106
Setting up the Mixed Platform Network	106
<i>Two Problems for Mixed Platform Networks</i>	107
<i>Sending and Translating Files</i>	108
For Mac-to-PC Connections	108
For Mac-to-NeXT Connections	108
<i>Large Graphic Files Degrade Network Performance</i>	109
Macintosh-to-PC	109
Macintosh-to-NeXT	111
Network Alternatives	112

Chapter Four: Groupware Implementations

Using the Right Tool for the Job	122
<i>When the Right Tool Isn't Always the Right Tool</i>	123
<i>Shared Space: A New Model for Collaboration</i>	125
Xerox Colab	126
Xerox Cognoter	127

Xerox Argnoter	128
Establishing and Connecting the Workgroup.....	130
Assessing Network Topologies.....	131
Assessing Cabling Media and Cabling Hardware	132
<i>Assessing Twisted-Pair LocalTalk Cabling Systems.....</i>	<i>136</i>
Apple's LocalTalk Cable System	136
Farallon Computing's PhoneNET System	136
Establishing Workgroup File Service	137
Dedicated File Servers	138
<i>AppleShare.....</i>	<i>138</i>
Nondedicated File Servers.....	141
<i>Personal Server Network</i>	<i>141</i>
<i>System 7.0 File Sharing</i>	<i>145</i>
Distributed File Servers	150
<i>DataClub</i>	<i>150</i>
Assessing File Service Options.....	155
Establishing Workgroup Printer Service.....	157
MultiFinder and System 7.0 Print Spooling	157
AppleShare Print Server	160
Establishing Electronic Mail Services.....	161
Workgroup Electronic Mail Considerations	162
QuickMail v2.2.3	164
<i>QuickMail Bridges</i>	<i>164</i>
<i>Remote Access</i>	<i>170</i>
<i>Bulletin Board</i>	<i>171</i>
<i>Groups</i>	<i>171</i>
<i>Mail Notification.....</i>	<i>172</i>
<i>Conferencing.....</i>	<i>174</i>
Microsoft Mail v2.0.....	176
<i>Server Administration.....</i>	<i>177</i>
<i>Mail Notification.....</i>	<i>178</i>
<i>Preferences</i>	<i>180</i>
Evaluating Electronic Mail Systems.....	181
Network Management.....	184
GraceLAN.....	187
NetUpdater.....	195
Teleconferencing	198
TeleFinder	199

Remote Network Access	209
<i>PhoneNET Liaison</i>	209
Virtual Meeting Facilitation	211
Timbuktu	212
Timbuktu/Remote	217
Group Authoring & Editing	219
MarkUp	220
Document Management	229

Chapter Five: Tutorials

Planning and Installing the Network	232
Plan Network	233
Run Cable	236
Establish File Sharing and Printer Service	238
Establish Electronic Mail Service	240
Test Network	241
Establish Screen-Sharing Service	245
Virtual Meeting Tutorial	246
Computer Conferencing Tutorial	250
Teleconferencing Tutorial	253
Project Management Tutorial	254
A Collaborative Session Using Inspiration with Timbuktu	261
Workgroup Data Backup Tutorial	269

Chapter Six: Creating the Collaborative Environment

Removing Inhibitions to Collaboration	278
Redefining the Meeting Ecology	280
The Role of the Facilitator	283
The Role of Group Memory	284
Arriving at Consensus	285
Benefits of the New Meeting Ecology	286
Managing the Information Flow	287
Using Satellites and Independents	290
Designing the Physical Space for Collaboration	290
Tenets for Success in the Collaborative Environment	291

Collaborative Problem Solving	294
Hypermedia Support is Crucial for Successful Collaboration	296
Interface and Design Issues	297

Chapter Seven: Collaborative Hypertext and Hypermedia

A Network of Nodes	304
Software Facility for Linking Nodes	304
<i>Link Directionality</i>	306
<i>Link Multiplicity</i>	306
Buttons for Arbitrary Links	307
Six Modes of Hypertext	307
Collaborative Hypertext Design Issues.....	309
Development Design Considerations	309
User Design Considerations	310

Chapter Eight: Problems With Groupware

Potential Interpersonal Conflicts.....	316
Workload vs. Perceived Benefit.....	317
Problems with Evaluating Tools.....	318
Hidden Agendas and the Gatekeeper Effect	318
Sacrificing Individuality to the Group Mind	319
Workgroup Homogeneity and The Groupthink Phenomenon	320
Broader Social Implications	321
Social Control Issues.....	322
What Good are Computers, Anyway?.....	324
The Mass Production and Merchandising of Information	327
Access Issues	328
Privacy and Security Issues.....	333
Intellectual Property Issues	336

Chapter Nine: Future Groupware

Virtual Reality	340
-----------------------	-----

Appendix A: Bibliography	343
Appendix B: Product Sources	363
Appendix C: Glossary.....	369
Index	405

Chapter One

Introduction and Overview

Groupware can be defined as a group of technologies, techniques, and services designed to help people collaborate more effectively, productively, and creatively. Groupware can consist of hardware, software, services, and support.

The model used in parts of this book is that of the meeting. Meetings can be held synchronously (all the members are together in time) or asynchronously (group members carry out discussions over a period of time and rarely, if ever, actually “meet” in real-time). A collaborative workgroup is comprised of at least two members and as many as can reasonably be supported. This can number into the tens of thousands in certain instances.

Vannevar Bush’s Memex

The name Vannevar Bush is not widely known within today’s computer industry. Most of the industry pundits believe that work done 50 years ago has little relevance today. It’s their loss, because Vannevar Bush’s ideas presaged a great deal of what we take for granted in computing today.

When he is credited at all, Vannevar Bush is usually credited with advancing the notion of non-linear writing, reading, and retrieval—the basis for what we now recognize as hypertext and hypermedia. Later in this book I will demonstrate

how hypertext and hypermedia technologies are extremely important aspects of collaborative computing.

Vannevar Bush made his most important contributions during the period of time at the end of World War II. In 1945 the United States was embroiled in the war and most of the country's efforts—including science and research—were focused on the singular goal of ending the war. In 1941 President Franklin D. Roosevelt named Vannevar Bush to supervise and coordinate all federally funded research through the Office of Scientific Research and Development.

In the summer of 1945, *The Atlantic Monthly* published Bush's seminal article, "As We May Think," a far-reaching piece that is credited—in hindsight, of course—as the impetus for a wealth of concepts central to the entire computer industry. It is a very chilling experience to read Bush's piece today, with the benefit of having seen so many of his visions realized, albeit in forms much more advanced than originally proposed. Bush accurately predicted high-resolution displays, fast information retrieval, and mass storage, all of which were as foreign to him as living on Mars is to us.

The central concept Bush proposed in "As We May Think" was a machine he called the "memex." The memex was a direct result of a need Bush perceived: how to manage the overwhelming amount of data and information that would be generated as the United States shifted its collective scientific focus from war-making to research and development.

One of the most daunting tasks confronting Bush in his effort to supervise and coordinate the nation's research efforts was the continual updating of technical information that was being generated by the scientific community at the close of the second world war. He became fascinated with the concept of microfilm and "As We May Think" was based on his ideas about it. The memex would be capable of holding all the writings generated in the scientific arena along with cross-references, links, and navigational trails among them.

At about the same time, Bush also proposed the Bush Rapid Selector, a microfilm tool that later was produced as a microfilm reader by Kodak and others. The key to the Bush Rapid Selector was various indices along the side of each microfilm strip. Bush's concept of associative "trails" were physical marks and sequencing cues embedded on the film. Of course we now recognize this concept as paths or tours in hypermedia parlance.

As World War II was winding down, and the scientists were free to pursue their individual interests, Vannevar Bush began to contemplate the direction and role of science in the future of society. Bush became more aware of the sheer bulk of research material prepared during the war effort and was continually confronted with the difficult process of searching across multiple documents.

Bush was convinced that the answer to his dilemma lay in the technology itself, as he noted in "As We May Think." "The world has arrived at an age of cheap complex devices of great reliability," he wrote, "and something is bound to come of it."

Microfilm was the era's state-of-the-art technology and it allowed scientific records to be compressed and stored in a minimal amount of space. Bush speculated that an entire encyclopedia could be compressed to fit on a single sheet of microfilm and that cheaply reproduced copies would provide access to vast amounts of information by virtually everyone that needed access to it.

Bush further speculated that there was a discernible difference between repetitive thought and creative thought and that the coming technology—including computers, voice-input devices, and scanners—would significantly enhance the creative process in individuals working together. Researchers would be freed from mundane repetitive mental tasks and have more time to spend on projects requiring creative thought.

Bush was well aware that the human mind operates by association and that, by extension, man would work best by applying these associative properties of thought. Most data are stored and sorted alphabetically and information is haphazardly found by hunting for it from subclass to subclass. Bush argued that an associative selection process could be mechanized, and that this process—while significantly slower in performance than the human mind—would possess the property of permanence rather than being of a transitory nature as are human thought processes.

Bush argued further that any specific bit of data would be accessible by entering a code, resulting in the display on screen of the document. Margin notes, annotation, and comments, according to Bush's vision, could be added at any point and associations could be constructed between two documents and displayed on adjacent screens.

The association, or "link"—complete with index—would be made *and stored* at the press of a button. The link could be easily recalled at any time and a new microfilm, consisting only of the linked pieces of text could be created and distributed separately from the original, underlying documents.

This mechanized system would allow for the creation of customized encyclopedias and other vast repositories of data, designed for a specific task or range of interests. Users could easily access the information via their individual memex machines and amplify the associations of other authors with their own notes and comments.

Vannevar Bush's tool for thought, the memex, was never built and his original vision is only now beginning to materialize.

Douglas Engelbart's Augmentation Work

Without a doubt, the basis for most of what we now recognize as groupware or collaborative computing is based on the ideas of Douglas C. Engelbart. Engelbart was one of the second-generation visionaries that followed Vannevar Bush's path, and the first to realize that while these new tools were going to revolutionize our access to information, some sort of framework was going to be needed to structure the capabilities we were going to be confronted with. His concept of "augmentation of the human intellect" grew from those concerns and have provided the framework for most of the modern computer industry.

Regarded by his contemporaries as something of a well-intentioned crackpot, Engelbart eventually received Department of Defense funding in the 1960s through the Advanced Research Projects Agency (ARPA). Ideas formulated at Engelbart's Augmentation Research Laboratory at the Stanford Research Institute (SRI) include the mouse, windows, electronic mail, and computer conferencing. Unfortunately, Doug Engelbart's augmentation system for the knowledge worker, however, remains to be implemented in a manner he considers to be acceptable.

Douglas Engelbart read Vannevar Bush's "As We May Think" piece while he was a radar technician stationed in the Philippines during World War II. The ideas proposed by Bush festered until Engelbart was 25 and living in the halcyon California of the 1950s. he decided to address in some manner the fact that the most pressing problems facing society were growing faster than the tools used to solve them. Engelbart envisioned a tool that would give a small workgroup, working collaboratively, a better chance at solving problems that were becoming ever more complex.

NLS and Augment

Engelbart understood that what was needed was not a new way to expand knowledge, but rather, new ways of discovering where to look for specific answers—answers that were already in cold storage somewhere. He also perceived a need for better communications tools between the individuals working collaboratively on complicated problems. Although Engelbart's augmentation system and attendant tools remain "in process," the underlying framework came to him in a flash. He recounted his experience in his groundbreaking 1962 article, "A Conceptual Framework for the Augmentation of Man's Intellect."

"When I first heard about computers, I understood from my radar experience during the war that if these machines can show you information on print-outs, they could show that information on a screen. When I saw the connection between a television-like screen, an information processor, and a medium for representing symbols to a person it all tumbled together in about half an hour. I went home and sketched a system in which computers would draw symbols on the screen and I could steer through different information spaces with knobs and levers and look at words and data and graphics in different ways. I imagined ways you could expand it to a theater-like environment where you could sit with colleagues and exchange information on many levels simultaneously. God! Think of how that would let you cut loose in solving problems.!"

Six years later, in 1968, a working system based on Engelbart's augmentation ideas was up and running. The system, called NLS (for oN Line System) included such advanced features as electronic mail, computer conferencing, multiple windows on screen, and a mouse. NLS was designed to allow anyone to read material written by anyone else and make comments and link to other documents from any terminal

connected to the system. The NLS system, in basically its original form, is still offered today as Augment by McDonnell-Douglas and is used mostly by the Air Force, although it is accessible via Tymshare. Douglas Engelbart has gone on to form the Bootstrap Institute in Palo Alto, Calif.

Augment, as the first collaborative computing system to actually succeed, includes features such as cross-referenced links between documents and the ability to expand or contract the information at the user's whim; tailoring the information flow to the individual user's needs. The amount of information available is further customizable via "viewing filters" which allow the user to specify the level of detail for a particular document.

The Knowledge Workshop

Concepts from both Augment and NLS made up what Engelbart and his contemporaries loosely referred to as the "Knowledge Workshop." Within the Workshop, any user could log into the system via any connected display terminal. Once he or she was logged in, all owned files as well as any files that were shared among the workgroup would immediately be accessible. Files could be read, new files could be created, shared files could be annotated. In addition, messages that were not connected to any document could be sent—immediately—to other Workshop users. Documents were easily transferred to other workgroup members simply by "releasing" them. No paper changed hands, and the transaction was perceptually immediate. Documents could be released to others for their comments and annotations, and the Workshop user would have common access to other workgroup members' shared documents.

If all this sounds vaguely similar to the modern Macintosh user connected to a local area network, it should. The basic concept of file servers is identical.

What separates Douglas Engelbart's "Knowledge Workshop" vision from current workgroup practice is the absence of paper and its attendant paper handling. Paper is eliminated at all levels. If you wrote—especially if you wrote a lot—this approach meant the end to lost notes scribbled several days earlier on napkins, match books, or other scraps of paper. Within the Knowledge Workshop, all of one's writings are available immediately, right *there*. Cross-references, footnotes, sidebars, and annotations are instant and painless. The Workshop promised an end to the time-consuming paper chase; looking for that scrap of paper containing last night's brilliant thoughts that just *has* to be here *somewhere*.

Another feature embodied within the Knowledge Workshop was the support for two people to work in a collaborative manner on the same document or set of documents. Two individuals, connected via telephone link, could work together on a common document: changes made by one person on her screen were immediately reflected on the other person's screen. No longer were geographically dispersed workgroup members subjected to the time delays of revision-by-mail. All revisions could take place in real time, or at least a reasonable facsimile of real time.

NLS and the Knowledge Workshop used a command language rather than the graphical interface familiar to Macintosh users. Although the language was fairly straightforward and the commands themselves were mnemonic (I for insert; M for move; D for delete; etc.), the user was expected to memorize the commands and they had to be entered in exactly the correct sequence. While such a system may have been acceptable to computer scientists, it was too much to ask of the "real" people Engelbart hoped to attract to the system. The system did, however, use the mouse as a pointing device to inform the system where the user was pointing on the screen.

Augmentation

Engelbart, again in his seminal "A Conceptual Framework for the Augmentation of Man's Intellect," defined augmentation—the term and concept—to mean, "increasing the capability of a man to approach a complex problem situation, gain comprehension to suit his particular needs, and to derive solutions to problems. Increased capability in this respect is taken to mean a mixture of the following: that comprehension can be gained more quickly; that better comprehension can be gained; that a useful degree of comprehension can be gained where previously the situation was too complex; that solutions can be produced more quickly; that better solutions can be produced; that solutions can be found where previously the human could find none."

Not only was Engelbart's intention to define and create new tools, but to define new ways of *working* with these new tools.

An appropriate example of augmentation is the concept of writing. Before human beings knew how to write they could only transmit ideas by telling each other. This oral tradition today survives in some cultures and even as parts of our own culture. Once humans learned how to write, they could communicate their ideas among themselves and have a permanent storage archive of their writings. Writing enabled the culture to become more informed by the sheer mathematics of the writer reaching more than one audience at a time. Computer screens take the tradition one step further. No longer confined to the printed word, ideas contained as light elements on a display screen promise to reach even vaster audiences and vastly enhance our individual "reachability" in both directions.

Central to Doug Engelbart's idea of augmentation of intellect was a redefinition of what we recognize as a concept. For Engelbart, a concept was something that, like thinking it-

self, evolved, and outmoded concepts could be readily replaced by other concepts. In addition, he felt that human thought processes and what he called “concept structures” could not only be monitored and studied, but amplified as well. To quote again from his original paper, “We view a concept to be a tool that can be grasped and used by the mental mechanisms, that can be composed, interpreted, and used by the natural mental substances and processes. The grasping and processing done by these mechanisms can often be accomplished more easily if the concept is explicitly represented by a symbol.”

This realization—that the human is aided in the grasping of concepts if the concept is represented by a symbol—led directly to the concept of a hand-held tool used as a pointing device for manipulating representative symbols on a computer screen: what we recognize today as the mouse. Engelbart went on to explain that a concept structure most often evolved on a cultural basis, either on a wide-spread or individual basis. He acknowledged that it was also—although with less frequency—something that could be “designed or modified” and that through appropriate modifications, these structures would improve the individual’s ability to understand the most complex problems confronting him and subsequently his ability to reach more insightful solutions to these most pressing problems.

A Conceptual Framework and Augmentation Means

The “conceptual framework” that Engelbart based his work on was in itself designed to be a specific plan for his own augmentation research. He quickly discovered that the basic principles applied to both the individual and the societal levels of experience. Engelbart proposed that by designing appropriate systems that would work in accord with human thought processes—if we could design systems that worked the way people worked—a synergy would result. Fully aware that the human mind is capable of only small steps, and that

each successive step relies on and builds upon previous steps, Engelbart felt that the resulting synergy was not capable of producing any larger steps, only more sure-footed ones.

Engelbart referred to the extension of human capabilities within his system as “augmentation means.” He further defined the augmentation means into groups of four basic classes:

- The *artifact* class of augmentation means referred to the human capability of manipulating symbols and physical objects to make themselves more comfortable.
- The *language* class addressed the manner in which the human mind organizes his world view into the concepts that the mind uses to create a model of the world and the symbols that are attached to those concepts in the thinking process.
- The *methodology* class spoke directly to the procedures employed by the individual in any problem-solving exercise.
- The *training* class of augmentation means were the conditioning needed to make the other three augmentation means work effectively.

Based on his idea of augmentation means, Engelbart observed that the augmentation means served to break up large, complex problems into more manageable chunks, allowing the individual to approach the problem as a series of small steps. He called the structure of the small steps “process hierarchies.” Although he recognized that each small step—each sub-process—was itself a process, Engelbart also realized that the human being never uses a “completely unique process every time he performs a new task.” We don’t reinvent the wheel each time we are confronted with a new problem; instead we build upon what we already know using what we already know.

To Engelbart, it was clear that there is a finite number of “tools” with which to fashion new solutions. That finiteness, however, in no way impacted upon the solutions available. As one of the heroes of the 1960s, Mr. Natural, opined, we have to use the right tool for the job. Even with a finite number of tools at our collective disposal, few of us ever become proficient with more than a handful of them. We continue to reuse tools that have worked in the past when we confront new problems. The down-side of this, of course, is that many of us tend to see every problem as a nail if the only tool we’re proficient with is a hammer. Engelbart envisioned a way to avoid that “tool-bias” limitation.

The key to Engelbart’s vision of bypassing our built-in tool- and process-biases was what amounted to screen-based text editing: what we now know as word processing. In an era when it cost hundreds of thousands of 1960s dollars to produce a computer capable of on-screen text editing, Engelbart was correctly predicting that such appliances would become commonplace within his lifetime.

Ted Nelson’s Collaborative Computing Work

Theodor H. Nelson is most widely known as one of the foremost researchers in the field of hypertext and hypermedia and as the coiner of both terms. Hypertext and hypermedia are important subsets of groupware technology, and most of Nelson’s work crosses over into the collaborative computing arena.

Xanadu

Nelson’s pet project of the last 30 years—the Xanadu Project—is only now beginning to show signs of emerging from the vapor. Xanadu, as conceptualized by Nelson and implemented by the Xanadu Operating Company, is a glo-

bal information repository and network. Ted Nelson refers to this repository as the "magic place of literary memory." Xanadu will eventually be comprised of many thousands of nodes throughout the world, some of which will exist as fast-food-franchise-like establishments Nelson refers to as "Silver Stands." When Xanadu is realized, many thousands of users will have simultaneous access to mountains of information. Of course, Nelson himself acknowledges that the name Xanadu is based on Coleridge's unfinished poem "Kubla Khan," so there are no guarantees.

The Xanadu system was originally conceived as a solution to a very real problem Nelson perceived while he was a student at Harvard. He needed a note-keeping program that would serve as a repository for his thoughts. Xanadu, however, quickly grew in scope to encompass idea creation, thought organization, backtracking, alternate versions, automatic cross references, text manipulation, and a complete electronic publishing system including an automatic royalty-logging mechanism.

Within Xanadu, every document in the repository has links to what it was drawn from and to those documents that draw upon it. The linkages are an electronic form of footnotes that appear in a separate window.

Nelson describes a unique storage system for Xanadu, which he called "xanalogical storage." The xanalogical storage model is based on a single repository that is shared across the system itself and at the same time organized in a myriad of different ways. Originally designed as a vast dumping and storage ground for textual information, the xanalogical storage model quickly grew to encompass all interactive media as well and began to be seen more along the lines of a heritage preserve accessible to all community members. Since the knowledge base comprising Xanadu is a single shared pool, the contained materials can be continually rearranged to meet the needs of varied individuals without losing any of the existing organization structures. Sounding like a huge

software program, in actuality, Xanadu is relatively small and relies on the cooperative processing power of all the nodes on the network. This architecture allows the system to be run by various individuals, under shared localized ownerships, rather than being subjected to the potential tyrannical rule by a single body.

Publishing on the Xanadu system will revolve around a unique royalty structure. All users will have access to all public documents. Each time a user accesses a given document, a royalty, albeit a minuscule royalty, will be returned to the originating document. The royalty would be a proportion of what Nelson calls the "byte-delivery charges," which amount to between ten and twenty percent.

Material published as a "public" document on the Xanadu system may not be removed from the system except by a six month notice or court order. The author or publisher who wants his or her material to be available to everyone but also wants to reserve the right to withdraw the material at any time can simply publish the material as a private document with unrestricted distribution. Additionally, any user can publish "collaged" or "windowed" documents that have a variety of degrees of "ownership." These documents would have their origins in public documents that are freely accessible to all. While no user would be restricted from any public document, private documents would have access restrictions specified by the author or owner of the document. Private documents would be accessible only to the owner of the document and whomever has been granted access privileges.

According to the Xanadu design specification, no limitation or restriction can be placed on the use of materials found within the system, largely because there is no practical way to implement such a restriction system. For example, if someone wanted to alter one of your documents—completely changing your point through annotation—he or she would be completely free to do so. The beauty of Xanadu,

however, is that the unadulterated original is always readily available, no further away than a link or three.

Since the user has paid a royalty fee to receive the material, he or she is free to do most anything with the information within the realm of existing copyright law. While any published material on the system may be printed out or copied to disk by a user, any link connections are lost and the user is left with what Ted Nelson describes as an "inert, non-interactive copy.... Remember the analogy between text and water. Water flows freely, ice does not. The free-flowing, live documents on the network are subject to constant new use and linkage, and those new links continually become interactively available. Any detached copy someone keeps is frozen and dead, lacking access to the new linkage (and, if there were any substantial body of in-links at the time the copy was made, probably most of those as well)."

While inside the Xanadu system, the user has much more freedom with the published material he or she finds. Once the material leaves the system, however, normal copyright law would apply. Therefore, any use outside of fair use would be as illegal as any other form of copyright infringement. Since all material on the system would be time and date stamped, the risk a publisher runs by making his material public (as opposed to private) would be relatively small.

The Xanadu design specification also allows for private sale of materials stored on the system for those not satisfied with the standard royalty structure. In the implementation envisioned by Nelson, anyone may store their material as private documents and sell access rights to others.

As complicated as Xanadu sounds, it's imperative to understand that the system is nothing more than a storage manager. What Nelson is striving for is release from the confines of what he calls the "tyranny of the file." Xanadu is built upon a single storage container that can not only be shared between virtually any number of users but also simultaneously organized in a virtually unlimited number of ways.

The goal is to make it easy for any user to expand his thought processes by physically creating new material from bits of pre-existing, thought-building blocks.

Ted Nelson sees the current state of networking and electronic publishing as a hopelessly embroiled group of small, mutually resentful bodies. He believes that a universal hypertext network would change that by supplying stored text and graphics on demand from anywhere to anywhere else on the network. Such a situation would, according to Nelson, render information "an elemental commodity, like water, telephone service, radio, and television." More importantly, Nelson holds that such a development would change the basic structure of information "with all its intrinsic complexity and controversy, and provide a universal archival standard worthy of our heritage and freedom and pluralism."

Thinkertoys and Super Virtualities

Ted Nelson finds himself in the company of Doug Engelbart and Vannevar Bush in maintaining that the greatest problems facing society involve "thinking and the visualization of complexity." Where Engelbart proposed a set of tools that were useful in augmenting the human intellect, Nelson proposes a set of tools that will enable us to envision complex alternatives to a given problem.

Nelson refers to these metatools as thinkertoys and defines the crux of any thinkertoy as the ability of the device to allow things of varying levels of complexity to be inter-compared and subsequently inter-comprehended via the interconnections of the devices themselves. Ted Nelson is very specific in providing anecdotal evidence for implementing such tools: assessing alternative designs and theories, comparing successive drafts of documents, and identifying discrepancies in courtroom testimony. The underlying premise of the thinkertoy is that although the in-

terconnections between vastly different problems appear to be themselves vastly different, in actuality, they are more similar than dissimilar.

Significant differences between types of problems, remain, however, leading Nelson to propose the most general of approaches to problem-solving. One of the techniques he specifies is something he refers to as "collateration," the "linking of materials into 'collateral structures.'" Structures could be said to be collateral, according to Nelson, if there are specific links between them, although the specific sequences of the connections may be different.

Ted Nelson's concepts are quite difficult to grasp in the abstract, but each of his ideas is built on the notion of simplicity. For Nelson, the guiding principal of any computer system, regardless of what it is designed to do, is ease of use. In one of his most important books, *Computer Lib/Dream Machines*, Nelson states his premise quite clearly: "If it is desirable that computer systems for *simple-minded* purposes be easy to use, it is *absolutely necessary* that computer systems for *complicated purposes* be simple to use." Many would criticize such a position as overly simplistic and patently obvious, but such tools are notoriously difficult to implement. Power and apparent simplicity are not mutually exclusive in the view of Ted Nelson and his greatest acumen may very well be his gift of demystifying the technology underlying the implementation of these important tools.

Nelson insists on apparent simplicity to such an extent that he admonishes systems as almost useless that are more complicated than what he calls a "ten minute system." A ten minute system, according to Ted Nelson, is one that can be learned by a novice and put to useful application in less than ten minutes. "I believe that interaction with computers can be at least ten times easier, Nelson states in *Computer Lib/Dream Machines*, "ten times more powerful, ten times more vivid; and that these are issues not of hardware but of virtuality design."

Chapter Two

Driving Forces of Groupware

The most important development of our lifetime has been the replacement of the second-wave, industrial economy with one based on information and knowledge. Instead of creating wealth with muscle, as we did in the industrial economy, we now create wealth by manipulating symbols—by converting information into knowledge.

The Information Economy

Most of us have held the belief that information and the knowledge distilled from that information are neutral. This was a mistake, because information and knowledge are both based on underlying facts and assumptions. The facts had to come from somewhere, and the assumptions are loaded with a variety of biases. Alvin Toffler articulately points this out in *Powershift*: "Virtually every 'fact' used in business, political life, and everyday human relations is derived from other 'facts' or assumptions that have been shaped, deliberately or not, by the preexisting power structure. Every 'fact' thus has a power-history and what might be called a power-future—an impact, large or small, on the future distribution of power."

The agricultural and industrial economies were based on a set of hierarchies and control structures. Power was controlled by whomever controlled the physical resources and influence was wielded largely through secrecy.

In an information-rich environment, the basic notion of *control* changes. Secrecy is abandoned simply because secrets are too hard—or too expensive—to keep. Openness becomes an economic imperative rather than a preference and many businesses are already becoming aware of this fact. Many experienced collaborators wrongly insist that a policy of openness helps stop the negative and at the same time helps prevent the positive unless a consensus has been arrived at privately. Sometimes, as you will see in later chapters, this is unavoidable but whenever possible such actions should take place in as open a manner as possible. Openness has been said to cause a loss of candor, less freedom of speech, distortion through simplification, an underlying concern about appearances, and general inhibition. I hope to show that these symptoms are not a result of openness but rather result from a pseudo-openness that is in fact overtly secretive.

In previous economies based on agriculture and industrial activity, various hierarchies of class were assigned based on ownership. In an information-based society, the concept of *ownership* changes drastically. When an idea or other piece of information is shared it is intangible and not subject to an exchange-based economy. The creator doesn't lose possession of the information when he shares it with the recipient. Even if the information is paid for, what's really being purchased is the *presentation and delivery* of the information, not the information itself.

Because information can only be shared, rather than exchanged, the power structure will undergo significant change. Since information can be shared more equitably and is virtually infinite, benefits will spread across the populace rather than contribute to the further concentration of wealth. Further, the restrictions on exchange-based infor-

mation transactions will lead to a greater level of diversity within the culture.

In agricultural and industrial economies, poverty was justified by shortages. Information is plentiful and education is the key to equitability. Access to information will become one of the most politicized issues in the coming years. Japan, for example, has already adapted its culture to this fact. Faced with an environment lacking in virtually all natural resources and a limited amount of tillable soil, the country bet heavily on its single greatest asset: it's own people.

In the previous agriculture and industrial economies, political power was wielded based on geography. In an economy based on intangibles, relationships will become more important than geographic proximity. Pyramidal organization charts—the kind enjoyed by most corporations—are distortions of reality. The people you work with on a daily basis don't work for you and you don't work for them. In effective businesses, you work together, even if the corporate organization chart doesn't reflect it.

A New Hierarchy

According to Alvin Toffler, "Any data, information, or knowledge that is communicated requires (1) a source or sender; (2) a set of channels or media through which the message flows; (3) a receiver; and of course (4) a message. Power players intervene at each of these points.... The identity of the Sender is, in fact, a crucial part of any message. Among other things, it helps us decide how much credence to give the message."

The set of hierarchies and control structures that worked in the agricultural and industrial economies won't work in the information economy. In an economy based on information, each of these hierarchies must be reexamined. The key to this reexamination is the dichotomy pointed out by Stew-

art Brand: "Information wants to be free. Information wants to be expensive."

Information Expands as it is Used

If I give you a copy of this manuscript on disk, I still have my original master, yet your copy is indistinguishable from the master. This type of information, because it has value over time, expands as it is used. Some types of information—usually time-critical information like stock market reports, weather forecasts, or a horse-racing tip—have little long-term value and therefore contract over time.

Information, by nature, is synergistic. The more we have the more we use and the more uses we put it to. Similarly, information is not scarce. Facts are abundant—probably never-ending—and we are so constantly bombarded with information that one of the crucial planning factors becomes uncertainty. The more we make our livings from manipulating information the more we will be confronted with the necessity of reducing information loads.

There are two limits to the growth of knowledge.

- *Time.* The time we need for interpreting, analyzing, and integrating information.
- *Capacity.* Our capacity to think in an integrated manner.

Groupware technologies will enhance the capacity limiting factor but will do little for the time factor.

Information is Compressible

Information is the most useful when it is compressed or concentrated. Compressed—or integrated—information is not

resource hungry. Relative to manufacturing technologies and processes, the creation and manipulation of information is downright stingy in the use of energy and other physical resources.

Information is Transportable

Information is more transportable than just about anything else we can imagine, because it's not tangible. This intangible quality of information is what allows virtual communities to exist free of geographic constraints. As this becomes more and more accepted in the business community, where people locate will become their sole decision. No longer will they be forced to relocate at the whim of someone else.

Information Wants to Leak

It's exceptionally difficult to contain information. It's like mercury and tends to be highly malleable in some ways, and very difficult to shape in others. Therefore, it's extremely difficult to monopolize information.

Information Wants to be Shared

In economic terms, information cannot be submitted to exchange transactions. Information wants a sharing transaction. If I sell you my computer, you have it and I don't. But if I sell you the information contained within my computer, we both have it.

The Collapse of the Information Float

Just as the speed of financial transactions has tended to eliminate the “float” many banks have counted on as a hidden source of profit, the “information float” has similarly been collapsed by a drastic increase in the speed of computers and information networks. Alvin Toffler, in his book *Powershift*, clearly identifies this trend when he speaks of increased transaction speeds. “New computers and communications networks not only make possible the variation and customization of existing products, and the invention of new ones,” he says, “but also drive transaction speeds toward instantaneity.”

Capital itself, or at least our concept of it, is dematerializing. Money is changing from tangible paper and coins to intangible symbols on a computer screen travelling quickly across a vast network. Money has become intelligent, and its very nature has changed right before our eyes. Alvin Toffler, again in *Powershift*, explains:

“Capital—by which we mean wealth put to work to increase production—changes in parallel with money, and both take on new forms each time society undergoes a major transformation. As they do so, their knowledge content changes. Thus agricultural-era money, consisting of metal (or some other commodity), had a knowledge content close to zero. Indeed, this First Wave money was not only tangible and durable, it was also pre-literate—in the sense that its value depended on its weight, not on the words imprinted on it. Today’s Second Wave money consists of printed paper with or without commodity backing. What’s printed on the paper matters. The money is symbolic but still tangible. This form of money comes along with mass literacy. Third Wave money increasingly consists of electronic pulses. It is evanescent ... instantly transferred ...

monitored on the video screen ... Third wave money is information—the basis of knowledge.”

Stanley Davis's Four Earmarks of the Information Economy

Stanley Davis, in his 1987 *Future Perfect*, points out the four earmarks of the information economy:

- any time
- any place
- no-matter
- mass customization

As we make the transition from an economy based on hard goods to one based on intangibles, the greatest intangible of all—time—becomes more and more important. Davis goes on to identify three underlying principles at work in viewing time as a resource in an information economy.

- “Consumers need products and services *any time* (i.e., in their time frame, not the providers’)”
- “Producers who deliver their products and services in *real-time*, relative to their competitors, will have a decided advantage.”
- “Operating in real-time means no *lag-time* between identification and fulfillment of the need.”

Information and the “Time is Money” Dichotomy

In the industrial economy, the phrase “time is money” implied that money is the most important resource; a resource not to be squandered and the yardstick used to measure the effective use of the resource. In the information economy, based on intangibles, the same phrase takes on an entirely

different meaning, and time becomes the resource and money the measuring device used to gauge its effective use.

Groupware technologies allow us to provide each other with access to mass-customized information and knowledge in real-time. We are enabled to put all our available resources to their highest and best use.

In the information economy tangible products will begin to decrease in value, or rather, their value will be based on the level of the intangible value they add to the product. Consider the Macintosh itself as a prime example. The Macintosh is comprised of mostly spare parts from a variety of manufacturers. Apple Computer Inc. combines these parts with the real value—the intangible information in the machine's ROM—to make the Macintosh unique.

Economist, businessman, and author Paul Hawken suggests that the "single most important trend to understand is the changing ratio between mass and information in goods and services." Stanley Davis takes Hawken's point one step further by extrapolating it into an equation:

$$\text{VALUE OF A DELIVERABLE} = \frac{\text{INFORMATION}}{\text{MASS}}$$

Davis illustrates the equation by providing the example of steel and computer chips. Steel contains a lot of mass with relatively little information while a computer chip has a great deal of information relative to its mass.

Information is usually seen as an intermediary good—except by information providers—that is, information is seen as a tool useful in the creation of other goods and services. The value of information, in and of itself, will grow and the proportion of intangible goods increases in the society at large. Since intangibles exist only in time, and can't be inventoried, it may take us a while to catch on to this reality. Stanley Davis expounds: "Since services dominate the new economy, since they are performed, not produced, and since

they don't exist until they are performed, everything depends on executing the performance well. This means that, in implementing strategies, delivery systems play a much larger role in the new economy than they did during the industrial period."

Claude Shannon's Information Theory

In 1948, Claude Shannon, a mathematician working at Bell Labs was assigned the task of figuring out how to squeeze more phone calls into each telephone line. Being a mathematician, he reached for a tool readily at hand and applied the Second Law of Thermodynamics (energy deteriorates over time) to the problem.

The result of Shannon's studies is what we have come to know as information theory; a theory that has proven to be much more widely applicable than the simple multiplexing of telephone calls.

Shannon demonstrated that no matter how well constructed a pair of telephone wires is and no matter what else you do to them (shielding, filtering, etc.) information would be lost and noise would be gained when you sent information across them. He called this process entropy.

Information theory holds, quite simply, that information tends toward entropy (disorder and disorganization).

Redundancy

Shannon offered redundancy as a solution to information entropy. Redundancy is basically the same information presented or transmitted repeatedly, perhaps in a slightly different way each time.

Redundancy is what makes the Macintosh interface work, even if Apple insists on calling it consistency. Redundancy is also what will make collaborative computing tools work, as Bruce Tognazzini, one of Apple's human interface designers points out in his excellent article "Principles to Design By—Part Two" in the October 1990 issue of Apple's newsletter for developers, *Apple Direct*:

"Within a single medium, redundancy takes the form of rules. Rules, by their very nature, limit the range of possibilities. Limits make interpretation of information easier. For example, let's say you've asked someone to give you a number between 1 and 10 and you think you hear that someone say, "sex." Since your possibilities are limited to digits, you can not only safely assume that the information was garbled in transmission, but you can also heal the error by finding the closest match, the number 6.... Numerous studies have demonstrated that around 75 percent of the information in the English language is redundant. Same with German, French, Italian, Japanese, and Chinese. Even primitive languages, such as Hawaiian, are highly redundant.... In fact, every human language has 75-percent redundancy.... Transmission of information is rife with possibilities for error. The transmitter can missend the message, the medium can add so much noise that the message becomes garbled, or the receiver may hear or even interpret it wrongly. Only through redundancy do we readers and listeners have a chance to heal, or correct, messages.... The more complexity a message involves, the more redundancy is required to heal it.... This principle predicts that an interface such as that of the Macintosh, rich in redundancy, should enable people to perform far-more-complex tasks than could be accomplished with the more abstract character-based interfaces of old...."

Groupware Technologies and the Mass-Customization of Information

It's difficult to address the paradox of mass-customization—especially the mass-customization of information inherent in the groupware technologies—because most of us still cling to an industrial-economy paradigm. The industrial economy paradigm is based on tangibles and economies of scale that are either undergoing significant transformation or no longer exist at all (if they ever did). Economies of scale dictated that a product could be *either* customized *or* mass-produced and that customized products have a higher unit cost. Using the economies of scale paradigm it is impossible to have both mass produced and customized goods; we must settle for one or the other. This is simply no longer true. The collaborative workgroups that are the best at mass-customizing their products will prove to be the most successful in the coming marketplace.

Stanley Davis employs the example of the floppy disk to illustrate this idea of mass-customization: "The world of mass customizing is a world of paradox with very practical implications. Whether we are dealing with a product, a service, a market, or an organization, each is understood to be both part (customized) and whole (mass) simultaneously.... The floppy itself is already a mass-produced commodity; the information any one floppy contains is what customizes it. And the customers have performed the final manufacturing and tailoring themselves."

Common sense dictates that a customer would much rather have a custom-made product than one that was mass produced if they were available for roughly the same price. Therefore, we can assume that there will be a broad market potential for the mass-customization of information and information by-products.

In an information economy, information is one of our most valuable resources. Groupware is about groups of people managing their collective information resources.

In 1920 less than 10 percent of the American population were information workers. Currently, more than half of the American populace are information workers. Forecasters estimate that by the year 2000, a full 70 percent of the American workforce will be information workers, 2 percent will work in agriculture, and 22 percent will be manufacturing workers. The French call this trend *l'informatisation de la société*, the informatization of society.

Information and the Anxiety it Causes

Our single biggest problem with information is the anxiety it causes. Richard Saul Wurman, best known as the author of *Information Anxiety* and various *ACCESS Guides*, has identified a condition he has termed "information anxiety" to refer to the state of increasing discontinuity between what we understand and what we think we should understand. We are being inundated with tremendous amounts of information, most of which is useless, some of which is important. Only recently have we realized that most of our information anxiety problems are a result of failing to recognize the interconnectedness of various pieces of knowledge.

It's important for us to realize the various components of what we commonly refer to as information; the raw material most of us will be working with in the near future. The task performed by most information workers is actually a four-phase refinement process that could be referred to only somewhat facetiously as the "getting of wisdom."

- *Data* can be recognized to be that body of unassimilated observations and objective facts that make up the raw materials information workers form into useful information.
- *Information* is nothing more (or less) than organized data. The data have usually been organized by someone else, however, and maintain a patina of objectivity.

- *Knowledge* is simply organized information. At this stage of the refinement process, the information is organized and internalized by the individual. Knowledge is information that you have managed to integrate with everything else you know.
- *Wisdom* is fully-integrated knowledge; bits of knowledge that have been made even more useful by the nature of their relationships to other bits of knowledge.

Our tools for turning the raw resource of information into knowledge haven't managed to keep up with the amount of information we are bombarded with on a daily basis. We grow anxious when we can't find what we need to know in a growing mountain of formless data. Anyone connected to a corporate electronic mail system who returns from a week's vacation to find 100 urgent messages has experienced information anxiety.

And the mountain is growing in geometric proportions. Imagine the repercussions of being inundated with more information on a daily basis—from a single information source—than our ancestors were in their entire lifetimes. Wurman states that “a weekday edition of *The New York Times* contains more information than the average person was likely to come across in a lifetime in seventeenth-century England.”

The question for producers of interactive media, then, is how do we overcome the information glut and make our materials compelling and accessible to users in their normal course of activity.

Richard Saul Wurman's Information Immediacy

Complex information is easily understood when it is presented in a series of hierarchical layers, like the layers of an onion. The recipient of the information should be able to peel away layer after layer, traveling deeper and deeper into

the specificity of the body of information. This information layering helps us refocus ourselves in relation to the continual bombardment of information.

Richard Saul Wurman offers a model for the appropriate layering of information. He calls his model the Five Rings, based on the degree of immediacy that various types of information have for us in our daily lives. According to Wurman, "the rings radiate out from the most personal information that is essential for our physical survival to the most abstract form of information that encompasses our personal myths, cultural development, and sociological perspective." A graphic representation of Richard Saul Wurman's model of the five rings of information immediacy is reproduced in Figure 2.1.

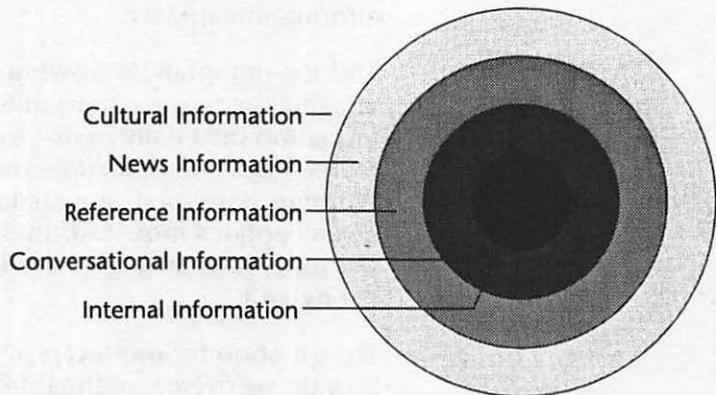


Figure 2.1 *Richard Saul Wurman's five rings of information immediacy.*

The outermost ring in Wurman's model is cultural information. This is our history, philosophy, and arts and represents our attempts to understand our culture. This is where information accumulated from the other layers is combined to build the information set that determines our attitudes and our beliefs.

The fourth ring is news information. Current events that have only a relatively minor impact on our daily lives comprise this level. Most of the information we receive from the media falls into this category. When we learn the name of the political candidates in Peru, for example, it's news information. This type of information, while not impacting directly on our lives, influences us in more subtle ways including the shaping of our individual conception of the world around us.

The third ring is reference information. This layer is made up of information sources including textbooks, encyclopedias, directories, and telephone books. We use reference information when we research specific information such as the year Portugal declared its independence or the phone number of a local city council member.

The second ring is conversational information. Conversational information is comprised of exchanges of information that we have with people around us. Wurman identifies this form of information as one of our main sources of information and the one we tend to ignore the easiest. Conversational information can range from the banal ("What's your name?" "What do you do for a living?") to the poignant ("What do you think about the proposed water referendum?"). This layer, however, is the most easily controllable. This is the layer Wurman tells us that we should focus most intensely on if we are to gain greater control over our own information spaces.

The innermost of the five rings is internal information, defined as that collection of cerebral messages that enables our bodies to function. When we feel pain or hunger, for example, we are dealing with internal information. This is the type of information that we have the least control over but that affects us the most.

The more closely you can integrate Wurman's five rings of information immediacy into your collaborations and groupware activities—especially those that are information-

based—the more accessible they will be for your associates and clients.

Workgroup Collaboration

Workgroup collaboration is the process of shared creation, or rather the process of creating value through shared creation. Michael Schrage, author of *Shared Minds: The New Technologies of Collaboration*, defines collaboration as “the process of *shared creation*: two or more individuals with complementary skills interacting to create a shared understanding that none had previously possessed or could have come to on their own. Collaboration creates a shared meaning about a process, a product, or an event.... But the true medium of collaboration is other people. Real innovation comes from this social matrix.”

Schrage goes on to define two types of collaboration, and acknowledges that we collaborate with the symbols people create in addition to our co-collaborators.

- *Formal collaboration* is the process we follow in meetings, task forces, and other highly-structured settings. This type of collaboration revolves around structures, procedures, and formal processes.
- *Informal collaboration* is identified by its casual and less-formal setting. Informal collaborations are likely to occur in hallways, at the coffee machine, or over lunch. This type of collaboration tends to be much less structured than a formal collaboration.

Both types of collaboration have common aspects. Good collaborators, regardless of the type of collaboration employed, tend to realize that the task they are confronted with is best addressed as a team rather than individually. Good collaborators recognize that to perform at their peak they need the input and feedback from peers and coworkers. In

Schrage's words, "they accept and respect the fact that other perspectives can add value to their own."

Workgroup collaboration is a means to any of three possible end results:

- Problem solving
- Creativity
- Discovery

Similarly, the means employed to arrive at the desired end result must operate within a set of five constraints.

- *Expertise.* The underlying reason for collaboration in the first place is that the task at hand is too great for a single person to accomplish.
- *Time.* Collaboration can take place either synchronously (in real-time) or asynchronously (in delayed- or compressed-time). The time constraint must be addressed appropriately, depending on the urgency of the situation at hand.
- *Money.* Budgets are a constraint in any situation.
- *Competition.* Market pressures must be addressed.
- *Conventional wisdom and common sense.* Society-wide prejudices and assumptions must be addressed. No one thought a four-minute-mile was possible until someone did it. Then, suddenly, a large group of people found themselves capable of it.

As much as you may have heard to the contrary, successful workgroup collaboration thrives on friction. Reasoned, even heated, argumentation is an asset not a liability to the collaborative process. Michael Schrage observes that the idea of argumentation "isn't pushy contrariness or a macho challenging of assumptions. It's really more a craft of explicitness. The arguments are a technique not just to flush out assumptions and underlying thoughts but to force the par-

ticipants to express their ideas so explicitly that they can be represented in a variety of ways, both tangible and intangible. The notion of 'winning' an argument is completely irrelevant—if not counterproductive—to the primary goal of clearly articulating ideas."

Workgroup collaboration has nothing whatever to do with the peculiarly American notion of "teamwork" and its attendant sports metaphors. Michael Schrage: "Collaboration is certainly not the bastardized American notion of teamwork and its Grantland Rice grab bag of tired sports metaphors—handing off, touching all the bases, striking out, and so on.... The issue isn't communication or teamwork—it's the creation of value. Collaboration describes a process of value creation that our traditional structures of communication and teamwork can't achieve."

Neither is effective workgroup collaboration a process of assembling the right pieces to a puzzle. The driving vision behind any collaboration is always dynamic; it never stagnates. Similarly, the workgroup members are constantly changing as well.

Michael Schrage has broken the collaborative process down into two distinct elements.

- *Conceptual collaboration.* Conceptual collaboration is the process of defining the overall goal of the collaboration and must precede the technical collaboration. Good examples of conceptual collaboration, according to Schrage, are a film director and a cinematographer blocking action and camera angles; Lennon and McCartney; and an art director and a copywriter discussing advertising themes.
- *Technical collaboration.* Technical collaboration is the actual attempt or attempts to solve the problem that was defined in the conceptual collaboration. Good examples of technical collaboration, according to Schrage, are James Watson and Francis Crick working

on the DNA model; Maxwell Perkins editing Thomas Wolfe's writing; and an architect and a contractor working together to arrive at the best solution for ventilating an office building.

Interpersonal Computing

In 1984, Apple Computer introduced its Macintosh computer; a computer designed to be as easy to use as an appliance. Apple stressed in its marketing that Macintosh was the computer for the rest of us. A large part of Apple's philosophy revolved around the notion of one computer for every person within an organization. Every Macintosh that has ever been sold comes with built-in networking capability—the foundation of groupware. Apple's marketing focus has subtly shifted from personal computing to interpersonal computing.

The 1990s will see the widespread development and integration of groupware; something that many pundits were sure would occur in the 1980s. Steve Jobs has gone so far as to identify the 1980s as the age of personal computing while the 1990s will be the age of interpersonal computing. According to *Fortune* magazine, "Software that will enable people to collaborate across barriers of time and space is one of today's hottest frontiers of computer research."

Because every Macintosh sold is equipped for local area network communications, Apple's Macintosh will be at the forefront of groupware, collaborative computing, and interpersonal computing. It's likely that Apple will revive its marketing prowess and attempt to lay claim to the concepts of groupware as it did four years ago with hypermedia. Apple has done a lot to further the notion of groupware and collaborative computing, but the concepts date back to the 1960s. Without doubt the first broadly successful groupware applications will appear on the Macintosh.

Interpersonal computing does not have to be all high-tech bells and whistles. A common low-tech empowerment example would be to attach a Macintosh to an overhead projector and use a scribe to record the group interactions and minutes of the meeting in real-time. This would form the basis of a simple Macintosh-based groupware system.

Michael Schrage acknowledges the effectiveness of the technique. Schrage and a collaborator were having difficulties expressing their ideas and points of view to each other. They scheduled a meeting with Bernie DeKoven, who facilitated meetings by a process he calls "technography" using a Macintosh connected to an overhead projector. Schrage's account of the meeting is powerful: "So Rob and I sat in Bernie's dimly lit office and talked about our idea. As we talked, our comments appeared on the large screen. We could see ourselves speak. Bernie technographed the discussion and we visually tracked where we agreed and where we diverged. We reacted both to each other and the screen. We defined our terms before our eyes. We could literally see ourselves skid off on tangents or become repetitious. At various points in the conversation, we took that jumble of on-screen comments and, using an outline processor, put them in some semblance of order. We ranked our priorities; we organized our thoughts; we literally mapped our discussion."

In Schrage's illustration, the technographer acted solely as a mirror; a function that could easily be served by the technology itself, given the proper tools.

Perhaps the most common experience we've all had using the Macintosh as an interpersonal computing tool is when we've worked with someone under tight deadline pressures. Imagine the following scenario and the tools that would be at your disposal. If you can't use the right tool for the job, use the tools at hand, but be careful.

Imagine you're a writer working diligently on a project with a designer and the project manager walks in and tells you that the project's deadline has been moved up. You've been

using LocalTalk to swap files with the designer and it's worked well. You send copy, she fits it in the design and sends the proof sheet back. With the accelerated deadline you no longer have the luxury of ping-ponging the copy and design work back and forth, so you copy your work on a floppy and head for the designer's office. Your files are transferred to her hard disk and she opens up the layout and begins to place the new material. Your copy doesn't fit. It's too long for the available space. In a panic you grab her keyboard and begin to edit your copy to fit the layout. Right there, in real time. She sees where she can free up a little more space without compromising the design and grabs the keyboard to make the changes. It works and you meet the deadline with time to spare by taking turns at the keyboard.

Imagine how much easier and more empowering the experience would be with tools designed or modified specifically for collaboration. This is the impetus for groupware and a prime example of why collaboration is an appropriate solution to many workaday situations.

Why Collaborate?

Computers aren't just personal anymore; they've become interpersonal and the days of using LocalTalk cabling as a long printer cable are over. Not only are computers a communications medium, they've become a medium for sharing communications and more importantly a medium for sharing understanding.

Collaboration is the natural extension of communication and six basic aspects of groupware technologies can be said to define an impetus for groupware:

- Conflicts are not two-sided.
- All problems are interdisciplinary.
- Solutions can be reached by widening the community.

- Consensus doesn't have to mean unanimous agreement.
- Voting doesn't work; consensus building does.
- Sometimes it's best to just agree rather than trying to agree on why you're agreeing.

History shows that collaboration is necessary for paradigm-shift-quality advancements in most, if not all, disciplines. Derek de Solla Price identifies this phenomenon as the "invisible college," the mostly informal and unstructured network we all form to share information. Examples of paradigm-shift-quality advancements that resulted directly from collaborations come easily to mind:

- James Watson and Francis Crick won the Nobel Prize in 1962 for their collaboration to discover the DNA double helix.
- John Lennon and Paul McCartney formed one of the most successful songwriting teams in history.
- Steve Jobs and Steve Wozniak collaborated to bring the original Apple computer to fruition.
- A very creative development team within Apple collaborated in the creation of the Macintosh.
- Bill Atkinson collaborated closely with Dan Winkler and other members of the original HyperCard development team to create the most innovative software program in recent memory. Atkinson acknowledges the collaboration in the first sentence of the Foreword to Dan Winkler's *HyperTalk 2.0: The Book* (which was itself a collaboration between Winkler and Scot Kamins). "I have fond memories of the year that Dan Winkler spent six days a week, working very long hours over at my house to create the original HyperTalk language."

Perhaps the most compelling reason for collaborating is that it's too expensive not to. Michael Schrage states that "the

cost of *not* collaborating is too high. Collaboration becomes an act of managing risk as well as managing innovation." Alvin Toffler offers even more evidence: "IBM alone connects 355,000 terminals around the world through a system called VNET, which in 1987 handled an estimated 5 trillion characters of data. By itself, a single part of that system—called PROFS—saved IBM the purchase of 7.5 million envelopes, and IBM estimates that without PROFS it would need nearly 40,000 additional employees to perform the same work."

The Collaborative Process

Collaboration has a lot to do with communication, but it's important to realize that the two are not the same. In fact, effective communication skills are important for effective collaboration, but in no way will the best communications skills guarantee a successful collaboration. Communication simply isn't enough. Michael Schrage notes, "Traditional modes of discourse in no way capture the subtleties, the bandwidth, the power, and the degrees of interaction necessary for effective collaboration.... The practical reality of collaboration is that it requires a higher order of involvement as well as a different approach to sharing and creating information.... The equation that *effective communications equals effective collaboration* is wrong.... Perhaps if you listen carefully and express yourself clearly, you will become an effective communicator. But the ability to communicate may have only a marginal impact on your ability to collaborate productively."

Irving Janis, in his 1982 book, *Groupthink*, acknowledged that the groups found running large organizations ping-pong their individual biases off each other through a mutual self-deception network. The result is a kind of management by, "I'll tolerate and support your biases if you'll support mine."

The largest barrier that communication poses for effective collaboration is the phenomenon that most effective communicators tend to become inseparably attached, or “married” to their communications. Collaborations centered around communications rarely are productive because the communicators tend to identify strongly with their individual contributions. Collaborations such as these quickly devolve into a set of individual expressions rather than a true collaboration. Individuals find themselves spending too much time advocating their positions rather than adding value to the group effort through perspective and comment.

The traditionally accepted idea that effective organizational communication—careful listening and clear speaking—results in understanding is simply wrong. This model is no longer appropriate because the level of complexity of the problems faced by today’s workgroups has drastically risen. Most organizations tend to try to address this flawed model by simply trying harder. “Caught up in the communication paradigm, they believe that more communication can compensate for their differences,” according to Schrage.

Schrage insists that “*what’s necessary isn’t more communication but rather a different quality of interaction—which is precisely what collaboration is.*” The different quality of interaction Schrage advocates is the creation of a shared space that allows the workgroup to manipulate ideas and information collectively. Most workgroup members, before exposure to the shared space paradigm, tend to treat each other as a sounding board and feedback device for self-expression rather than as collaborative partners. “In a communications-oriented environment, people discuss what they want to do and then go off and do what they think they’ve agreed upon;” says Schrage. “In a collaborative environment, people spend as much time understanding what they are doing as actually doing it. Vocabulary is defined precisely; imagery to illustrate ideas is agreed upon; individuals generate shared understandings that they couldn’t possibly have achieved on their own.”

Implications for Interpersonal Interaction

The collaborative process has serious implications for interpersonal interaction both inside of and outside of the collaborative environment. Michael Schrage has identified three dimensions of interpersonal interaction that change because of collaborative tools.

- *The role of language.* Collaborative tools enhance the visual nature of communication, but at the same time they require a new type of literacy. Schrage provides an apt illustration: "If the telephone renders mute someone who depends on body language to convey his real emotions, then collaborative tools might render the articulate egocentric mute as well.... Collaborative tools shift the role of spoken language away from transmitting ideas toward the construction of meaningful models."
- *The task of modeling processes and ideas.* Multimedia presentations have begun to raise the stakes of acceptable production values of presentations. The slickness without substance of these presentations has become tiring and boring to many, however, like eating too many doughnuts. Collaborative tools must be used to enhance the collection of new ideas and feedback and to encourage the participation of all workgroup members.
- *The perception of how others add value.* Schrage points out the possibility that collaborative tools will be used only to see other workgroup members as "entities to be mined solely for their knowledge and expertise." There is a great danger of seeing non-participatory workgroup members as non-contributors or even parasites.

The following chapters will address these potential problems with discussion and specific examples.

Chapter Three

Collaborative Computing With Groupware

In the future, computer networks will be as ubiquitous as the telephone system. Everyone will have access to everyone else and users won't have to worry about the network itself. We'll eventually put it to uses we haven't even dreamed of yet, but for the time being at least one person within every organization must be something of an electronic plumber and highway repairman.

Networking Basics

Networks provide five basic benefits to collaborative workgroup members:

- Resource sharing
- Information sharing
- Application sharing
- Accessing remote information
- Communicating

The Network as Electronic Highway

Networks, when well designed and implemented, serve to make differences and physical distances between computers appear invisible. The objective of any network is to allow individuals to have access and exchange information with others. From a day-to-day practical standpoint, however, most networks appear to be a rat's nest of cabling and connectors that magically get disconnected and a Pandora's box of software incompatibilities. It doesn't have to be this way, and this section will show you how to make your network as transparent as possible. It's all a matter of highway planning, road repair, and maintenance.

Networks are comprised of two elements, and each element has several sub-elements. Figure 3.1 illustrates a representative chart of the elements and sub-elements.

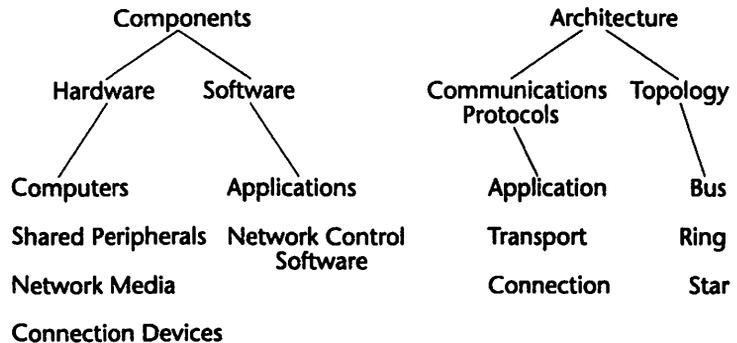


Figure 3.1 *Chart of elements and sub-elements of a network.*

Hardware Components

The hardware components of a network consist of four basic sub-elements:

- Computers
- Shared peripherals
- Network media
- Connection devices

Shared peripherals on a network can consist of any combination of printers, modems, and various servers. They are ancillary equipment that is shared by all members on the network.

The cables or other devices that carry the network's transmission signals are collectively referred to as the network media. Network media can be described in terms of three characteristics:

- Transmission speed
- Maximum length
- Shielding

Three types of cable are the most commonly used in computer network environments.

- *Twisted-pair wiring.* Twisted-pair wiring is popular because of its low-cost and ease of installation. It is the same type of wiring used for standard telephone installations and its characteristics make it ideal for low speed networks of moderate length. Twisted-pair wiring is comprised of two insulated wires twisted around each other so that external interference is balanced in the two wires. The two insulated wires are usually covered together with a pliable plastic casing. Figure 3.2 shows a cross-section of a twisted-pair cable.

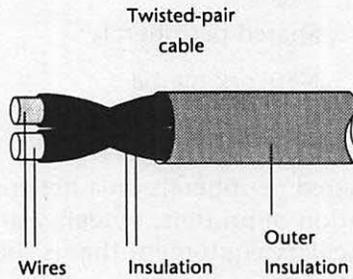


Figure 3.2 *Cross-section of a twisted-pair cable.*

- *Coaxial cable.* Coaxial cable is ideal for those installations that need a higher level of performance than can be accomplished with twisted-pair wiring. It is widely used in networks like Ethernet because of its high bandwidth and excellent shielding and is capable of speeds up to 16 million bits per second. Coaxial cable is the same wiring used for cable television and consists of a central wire surrounded with a three-ply sandwich of a conductive mesh between two layers of insulation. Figure 3.3 shows a cross-section of a coaxial cable.

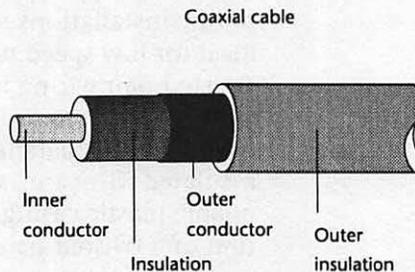


Figure 3.3 *Cross-section of a coaxial cable.*

- *Optical fiber.* Optical fiber uses light instead of electricity to transmit the network signal. A laser generates up to a billion pulses of light per second which is transmitted through very fine glass fibers. This medium is very resistant to interference, is capable of extremely fast speeds, and is best used for installations that require long-distance, high-throughput connections.

Each of the media types is appropriate for certain network installations. Appropriateness is determined by a combination of factors including network size, level and kind of use, and geographic proximity. Figure 3.4 is a table that illustrates the characteristics of the three popular network media options.

<i>Characteristics</i>	<i>Twisted-Pair</i>	<i>Coaxial Cable</i>	<i>Optical Fiber</i>
<i>Cost</i>	Low	Moderate	High
<i>Bandwidth</i>	Moderate	High	Very high
<i>Max. Length</i>	Moderate	High	High
<i>Immunity to Interference</i>	Low to Moderate	Moderate to High	Very high
<i>Ease of Connection</i>	Usually simple	Variable	Difficult
<i>Ease of Installation</i>	Variable	Variable	Difficult
<i>Reliability</i>	High	High	Very high

Figure 3.4 *Characteristics of three popular network media.*

Connection devices on the network refer to peripherals including Ethertalk cards and hardware routers.

Software Components

The software components of a network consist of the applications and network control software sub-elements. Applications software refers, in general, to network-specific applications such as screen-sharing software. Network control software refers to the file sharing software employed by the network as well as other software services such as electronic mail.

Network Architecture

The network architecture element is comprised of two sub-elements, communications protocols and topology, as shown in Figure 3.1.

Communications Protocols

Protocols are rules that govern how devices on a network interact with each other. Specifically, protocols oversee three functions:

- *Applications services.* Application services are the highest level network functions and enable an application to communicate with its equivalent on another computer connected to the network.
- *Transport services.* Transport services manage transmission control tasks and addressing.
- *Connection services.* Connection services are the lowest level network functions and control the physical transmission of information between nodes.

The ISO/OSI Model

Network protocols are referred to in terms of the *layers* of activity they perform. The Open Systems Interconnection (OSI) protocol layering model published by the International Standards Organization (ISO) is widely used to compare different network architectures. This model is referred to as the ISO/OSI model and is illustrated in Figure 3.5.

Layer	Name	Network Functions
7	Application	Providing the network services required by user applications. File service, etc.
6	Presentation	Converting the data formats as necessary for transmission by the network.
5	Session	Managing the sequence of interaction between two communicating devices on the network.
4	Transport	Assuring the reliability of a communication transaction between nodes on the network.
3	Network	Addressing and routing the data to the intended recipient network and node.
2	Data Link	Timing and coordinating the transmission as well as controlling access to the network medium itself.
1	Physical	Connecting to the network medium and generating the electromagnetic signal on the network.

Figure 3.5 *Layers and network functions of the ISO/OSI model.*

In the model shown in Figure 3.5, information moves from one node on the network to another by starting at the high-

est layer—the application layer—and moving through each of the lower layers until it reaches the physical layer. The actual transmission of the information takes place at the physical layer. The information travels up the layer hierarchy at the recipient node until it reaches the application layer.

Topology

Topology is the physical layout of the network and defines how each network component is arranged relative to other network components.

Bus Topology

A bus network employs a linear topology with each computer or other network device attached along the network cabling system from one end to the other. The network that uses a bus topology has two discrete ends instead of a continuous loop of cable.

Transmission packets on a bus are broadcast along the entire length of the cable, until they reach the recipient which accepts and reads the packet. Other nodes on the network simply ignore packets that aren't addressed to them.

An advantage of the bus topology is that no central control device is required. This allows computers to be added or removed from the network with little disruption and a single failure will not bring down the entire network. Figure 3.6 illustrates a typical bus topology.

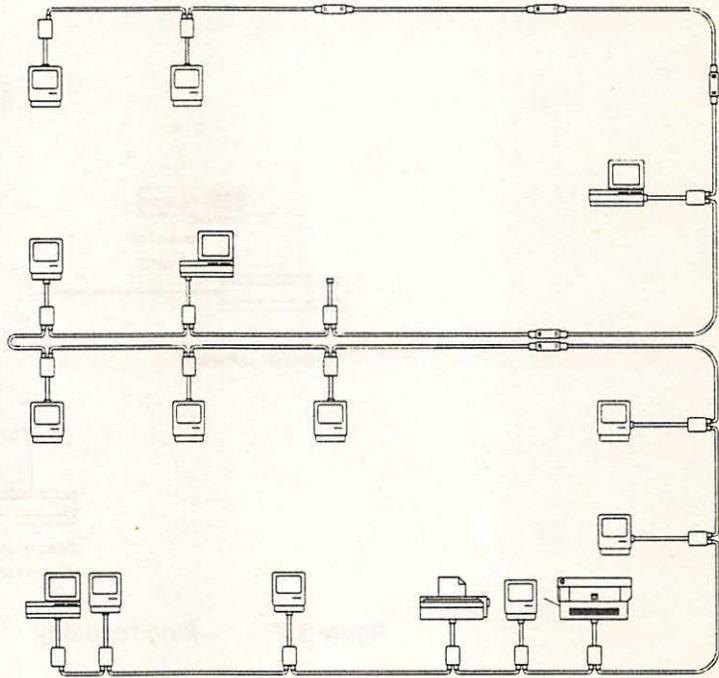


Figure 3.6 *Bus topology.*

Ring Topology

A ring network requires all network devices to be connected in a daisy chain, forming a closed loop. Transmission packets on the ring are passed from node to node until the recipient is reached. All data on the network is passed from node to node rather than being broadcast across the entire network. Figure 3.7 shows a typical ring topology comprised of both Macintoshes and IBM PCs.

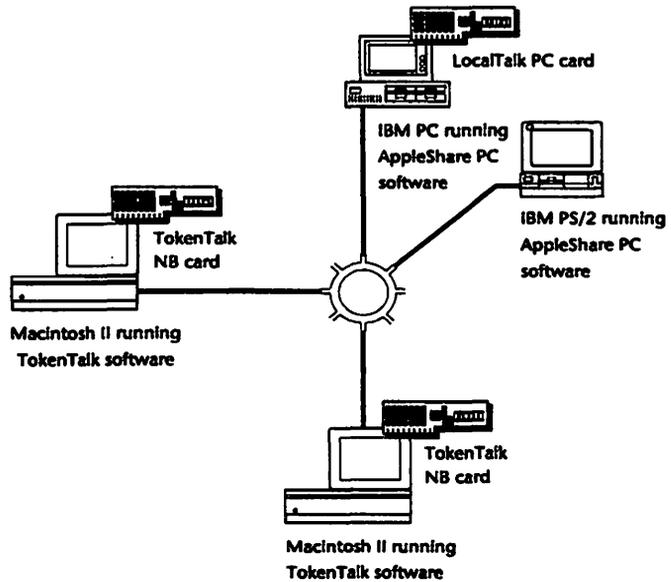


Figure 3.7 Ring topology.

Star Topology

In a star network, all nodes radiate from a central controlling node. The central controlling node can either be a computer running specialized software or a dedicated routing device, called a switch. The controlling node directs each transmission packet directly to the recipient node rather than broadcasting each packet across the network. Figure 3.8 shows a typical star topology.

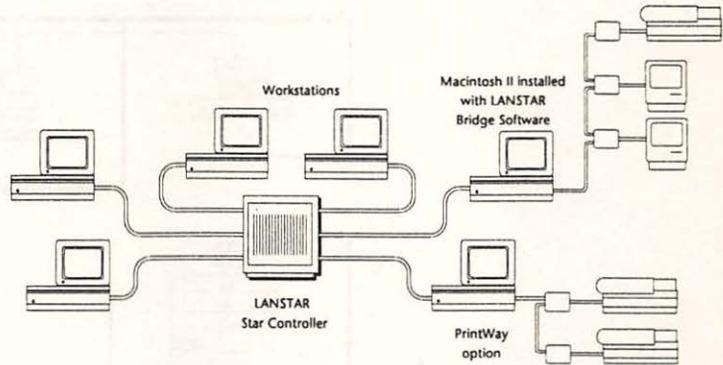


Figure 3.8 Star topology.

Categories of Networks

Any network falls into one of two categories. In addition, they may be linked together to form an internetwork.

- *Local Area Network (LAN)*. A local area network is comprised of any number of computers and shared peripherals which are physically connected together in a geographically centralized and limited area. LANs rarely span across multiple buildings, and the media is usually owned by the organization. A local area network's transmission speed is relatively high. A basic local area network is illustrated in Figure 3.9.

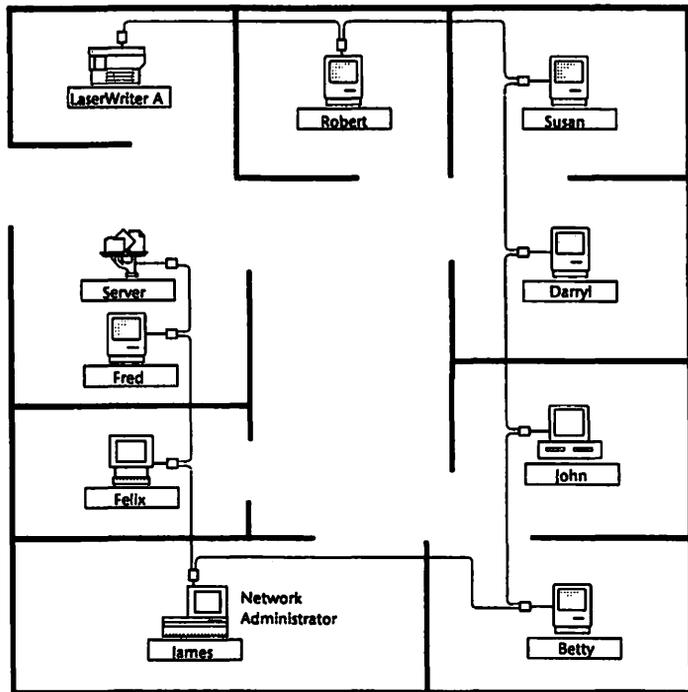


Figure 3.9 Basic local area network.

- **Wide Area Network (WAN).** A wide area network is comprised of any number of computers that are geographically dispersed and not physically connected together. WANs generally use telephone lines for their network connections and the media is usually not owned by the organization. A wide area network's transmission speed is usually significantly lower than that of a local area network. A wide area network, including several local area networks, is shown in Figure 3.10.

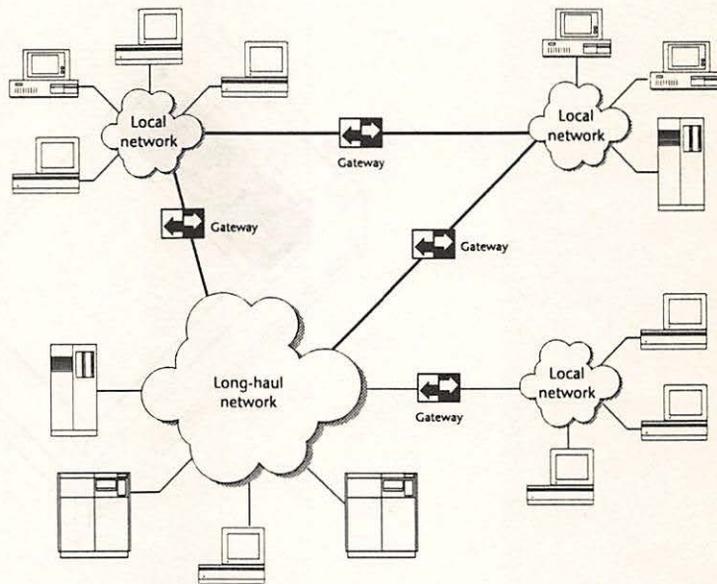


Figure 3.10 Wide area network including several local area networks.

Internetworks

When two or more networks are connected to each other they form an internetwork or internet. An internet can be any combination of local area networks and wide area networks. Four different kinds of devices can be used to connect networks, forming an internet:

- **Repeaters.** Repeaters are used to extend the physical network cable beyond its recommended maximum length or number of devices. By themselves they cannot be used to form an internet. An example of a common Macintosh repeater, Farallon Computing's PhoneNET Repeater, is shown in Figure 3.11.

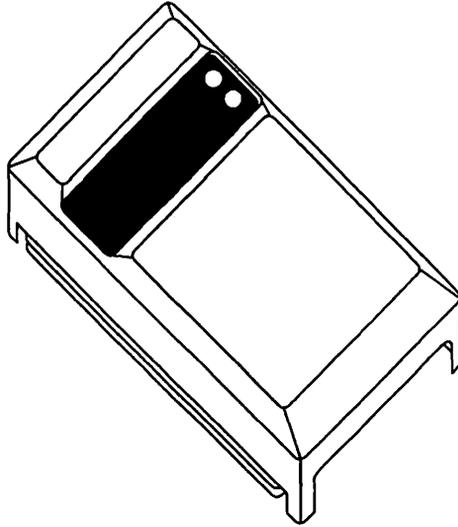


Figure 3.11 *Farallon PhoneNET Repeater.*

- **Routers.** Routers are used to allow connected networks on an internet to remain independent, each retaining its own identity and set of addresses. Routers select the most efficient path to a given recipient, allowing the network to operate faster and more efficiently. If a connection is broken along the path to the recipient, the router automatically detours to a different route. Common examples of network applications for routers are shown in Figure 3.12.

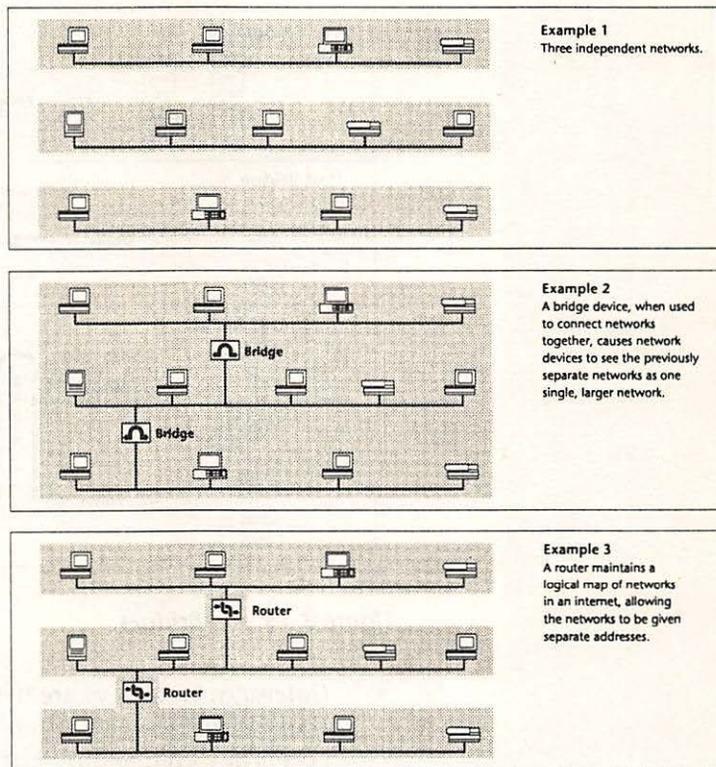


Figure 3.12 Routers.

- *Bridges.* Bridges use data link protocols to read a packet's node address and selectively filter or transmit the packet based on its destination. Networks that are "bridged" remain physically separate, but appear as a single network to other, external nodes. The three types of network bridges are illustrated in Figure 3.13.

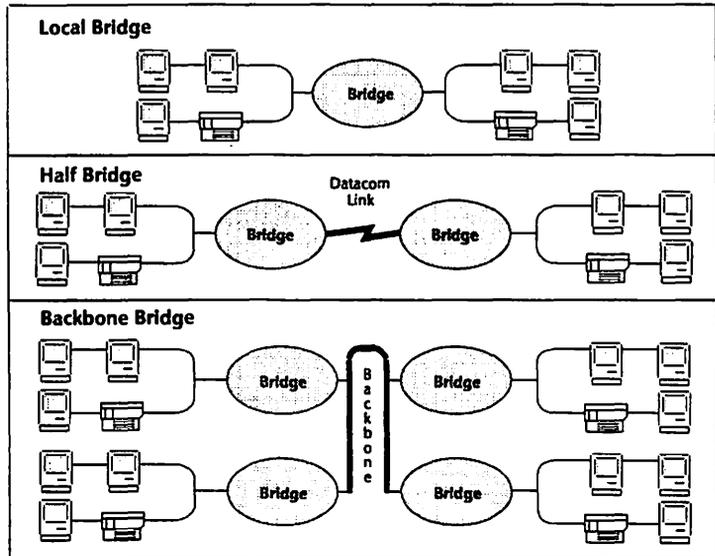


Figure 3.13 *Bridges.*

- **Gateways.** Gateways are best described as intelligent translators between different kinds of network protocols. They enable different kinds of networks to communicate. Figure 3.14 shows several of the most common network gateways used to translate disparate local area networks as well as a long-haul network.

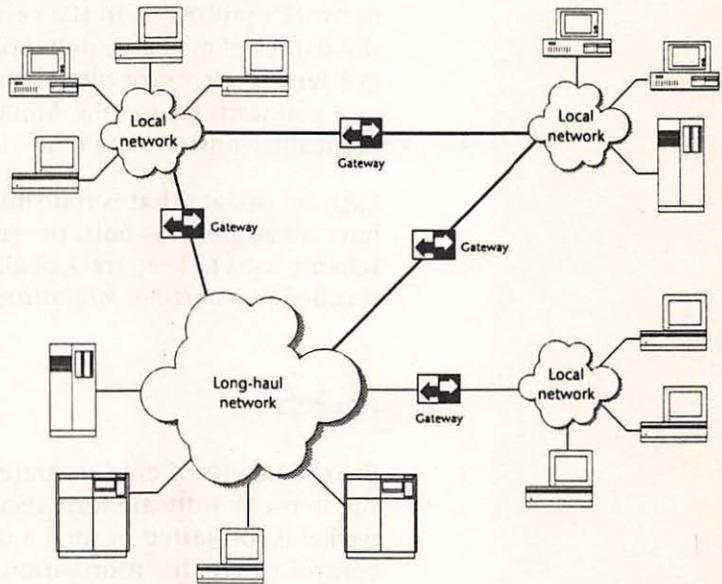


Figure 3.14 Gateways.

Nodes and Addressing

Networks are used to send and receive messages. These messages may be as simple as a two line electronic mail message or as complex as an interactive screen sharing session or series of file transfers, but when all the veneer is peeled away, network traffic consists of messages being passed back and forth. Every device connected to a network that is capable of sending or receiving messages is called a *node*. Computers, printers, servers, and other shared devices are all nodes.

Each node has an *address* which uniquely identifies it to other nodes on the network.

The way a network performs addressing is defined by the network's protocols. In the case of AppleTalk, you only see the names of available network services such as available printers, modems, or file servers. When you select the service you want to use, the AppleTalk network software automatically translates the name into the network address.

Each bit of data that is transmitted on the network has to have an address for both the sender and the recipient. The scheme used to keep track of all this addressing information is called *transmission formatting*.

Packets

Before the bits of data are transmitted across the network, the network software formats the data into *packets*. Each packet is formatted in such a way to allow the network to control where the information goes and what happens to it once it's received. A packet includes the actual information to be transmitted, the address of both the sender and recipient, and other information about the type of transmission that tells the receiving device what to do with the packet. The packet's format is determined by the host network's protocols.

Access Methods

Access methods are the ways in which a network internally manages access by multiple devices.

Apple's LocalTalk cabling system uses the CSMA, which is a *carrier sensing* access method. CSMA is an acronym that stands for *carrier sense, multiple access*. Using the CSMA access method, each device on the network checks the cable before beginning transmission. If the device senses a *carrier* (the signal on the cable) it assumes the cable is in use and

tries again later. This sounds like a long, time-consuming process but is actually happening many times per second.

Two devices could check the cable for a carrier at exactly the same time and begin transmitting. This would cause a collision between the two transmissions. A *collision detection* scheme is specified in the access protocol. Similarly, a *collision avoidance* scheme can also be used to retransmit information that is assumed to have collided on the network.

Token passing is another common access method used in some networks. Instead of giving each network node free access to the network, each node waits its turn for access to the network. A "token" or special sequence of bits is passed around from node to node allowing network access to the node that currently has the token.

Transmission Methods

Network transmission methods are the ways in which information is actually transmitted across the network.

The *broadband* transmission method transmits multiple signals simultaneously by dividing the medium into logical channels.

The *baseband* transmission transmits all signals simultaneously through a single channel. Most local area networks use the baseband transmission method.

A network's *bandwidth* is the range of transmission frequencies available and determines the potential throughput of the network. The higher the bandwidth, the higher the potential throughput.

Plugging In

Computers can be connected to a network in three ways. Additionally, any combination of these three methods can be used on a single computer.

Built-in Circuitry

Every Macintosh ever manufactured has built-in circuitry and software to support connection to a network. Apple's LocalTalk cabling system was originally designed for Macintosh computers and LaserWriter printers.

Add-on Circuit Board

Circuit boards for the Macintosh—such as Ethertalk cards—are available for network connections. Additionally, circuit boards are available for other computers enabling them to connect to LocalTalk cabling.

Modem

A modem is a hardware device that allows you to connect to a remote computer or network through an ordinary telephone line. The name modem comes from what the device actually does: it's a MODulator-DEModulator.

There are several standards issues that must be addressed in order for two modems to communicate with each other. In each instance, both modems must agree on the standard being used in order for the feature to be enabled.

The Consultative Committee on International Telephone and Telegraph (CCITT) is an international group comprised of technical experts responsible for developing international data communications standards. Organizationally, the

CCITT is part of the United Nations and its members include representatives from governments, common carriers, and modem manufacturers. The CCITT is responsible for establishing international standards for modem *modulation*, *error correction*, and *data compression*.

Modulation refers to the signaling method used by the modem and the two modems that are connected must use the same modulation method in order to communicate.

Negotiation is the procedure the two modems go through to determine which modulation method will be used for a connection. This is done by a series of tones sent by the modems to each other during the negotiation phase of the connection. Negotiation standards determine the sequence of actions that take place when a modem answers the phone.

Error Correction

Error correction standards enable modems to recognize errors during a transmission and to automatically request that the information be re-sent.

The V.42 standard is a CCITT specification that is similar to MNP Class 4, and in fact includes compatibility with all MNP standards through Class 4. V.42 provides a better protocol known as Link Access Procedure for Modems (LAPM). The LAPM protocol enables the modem to automatically re-transmit data that has been corrupted during transmission. This error correction assures that all information transferred during the connection will be free of any error.

Data Compression

Data compression standards govern a modem's ability to compress data before sending it, resulting in faster throughput. Depending on the kind of data that is being transmit-

ted, compression rates of up to 50% are common with text files, resulting in an effective doubling of throughput. In practice, however, most files are compressed before they are sent, and data compression within the modem is minimal.

The V.42bis standard is a CCITT specification that is similar to—but incompatible with—MNP Class 5. V.42bis offers approximately 35% better compression by compressing only that data that needs to be compressed, however, resulting in greater throughput. The protocol is intelligent enough to recognize when compression can be used and will disable it when it is not needed. V.42bis data compression requires the presence of V.42 error correction.

Microcom Networking Protocol (MNP)

The Microcom Networking Protocol (MNP) is a suite of error correction and data compression protocols developed by the modem manufacturer Microcom Inc. As MNP developed, additional standards were defined and added to the suite of protocols, although the most common implementations offer support of Classes 1 through 5.

- *MNP Class 1* provides less than 70% efficiency using an asynchronous, byte-oriented, half-duplex transmission. This error correction protocol is not used in modern modems.
- *MNP Class 2* is a full-duplex version of MNP Class 1 and offers less than 85% efficiency. This error correction protocol is not used in modern modems.
- *MNP Class 3* achieves just under 110% throughput by using a synchronous, bit-oriented, full-duplex method.
- *MNP Class 4* offers just under 115% throughput efficiency and is most effective on clean connections that travel minimal distance.

- *MNP Class 5* is a real-time data compression protocol and claims an increase of up to 50% in throughput efficiency. In practice, however, this figure varies widely depending on the type of data being sent. On files that have been compressed using Compactor or Stuffit, MNP Class 5 can actually decrease effective throughput.

Modulation Standards

Modulation standards dictate the speed, in bits-per-second, at which two modems connect.

Most 300 baud modems in North America adhere to the Bell 103 standard initially created by AT&T. Modems manufactured outside of North America, however, use the incompatible V.21 standard. Many modern modems can be configured to support either standard.

The Bell 212A standard is followed by most 1200 baud modems in North America. Modems manufactured outside of North America adhere to the incompatible V.22 standard. Most modems manufactured since the mid-1980s can communicate using either the 212A or the V.22 standard.

V.22bis is the international standard for 2400 baud communications. In addition to providing 2400 bps communications, the V.22bis standard also includes automatic baud rate detection and data rate fall back allowing connections at lower speeds that are common to both modems.

Standards for 9600 baud transmission rates have been in place for longer than most people realize. The V.29 standard governs 9600 baud half-duplex communications and is widely used in Group III facsimile (FAX) machines. The V.32 standard is the 9600 baud full-duplex standard and was prohibitively expensive to implement until the late 1980s. Because of the high cost associated with implementing V.32, several modem manufacturers resorted to proprietary trans-

missions for high speeds that were generally incompatible with modems from other manufacturers. These proprietary technologies included offerings from Hayes, the U.S. Robotics HST design, and the Telebit PEP method.

The V.32bis standard was scheduled for CCITT definition and approval during my writing cycle for this book (early 1991). V.32bis will allow modems to establish 14,000 bits-per-second, full-duplex connections.

Network Services

The network services available to users on the network are the specialized applications that they perceive to be the "network." These network services include applications that allow workgroup members to share information, share network resources, and communicate with other users.

File Service

One of the most common network services is the file server. A file server is specialized software that makes files and applications available to any member of the workgroup that is connected to the network. These files and applications can be copied from the file server to the user's local hard disk or they may be opened and used on the server as though they were physically located on the user's local hard disk.

Any file or application that is needed by various members of the workgroup are candidates for storage on a file server. In the case of a documentation department, for example, corporate logos, templates, style guides, boilerplate materials, and the like are all likely candidates for storage on the file server.

Additionally, any software that has been licensed for multiple users can also be placed on the file server. In theory, this

allows members of the workgroup to use the same application and document set simultaneously. In practice, this is sometimes more trouble than it's worth and various scenarios will be discussed later.

There are three types of file servers:

Dedicated file server

A dedicated file server requires a computer to be dedicated to the task of providing file service to network users. In the Macintosh world, Apple Computer's AppleShare is the most widely used dedicated file server and consists of a dedicated Macintosh, the AppleShare file server software, and one or more large capacity hard disks.

The main benefit of a dedicated file server is centralized administration and better performance in certain, high-load situations.

The main drawback of a dedicated file server is that a dedicated Macintosh is required and cannot be used by a workgroup member. In general, larger workgroups benefit most from a dedicated file server.

The most common dedicated file server for the Macintosh is AppleShare. Its structure is illustrated in Figure 3.15.

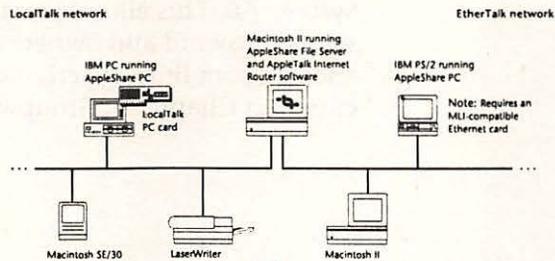


Figure 3.15 *AppleShare structure.*

Nondedicated file server

A nondedicated, file server is a software-based network service running on each workgroup member's computer. This approach requires each workgroup member to share the load of file service for the network.

The main advantage of a nondedicated file server is that a dedicated Macintosh is not required as the server machine, thereby reducing the costs for smaller workgroups.

The main drawbacks of a nondedicated file server are decentralized administration (which is seen by some as an advantage in and of itself) and reduced performance in certain, high-load situations.

A nondedicated file server allows users to designate specific files or folders on their own disks that can be accessed by others. This allows each user to have direct access to other users' files.

Apple's System Software 7.0 will include nondedicated file service when it is released. Until then, the two most popular nondedicated file servers are TOPS' Sitka and Interactive Presentation Technology's Personal Server Network.

The nondedicated file service built-in to System 7.0 is illustrated in Figures 3.16 and 3.17.

Figure 3.16 shows the network Sharing Setup dialog box for System 7.0. This allows you to identify your computer, assign a password and owner name, and enable file sharing and program linking activities. These features will be discussed in Chapter 4: Groupware Implementations.

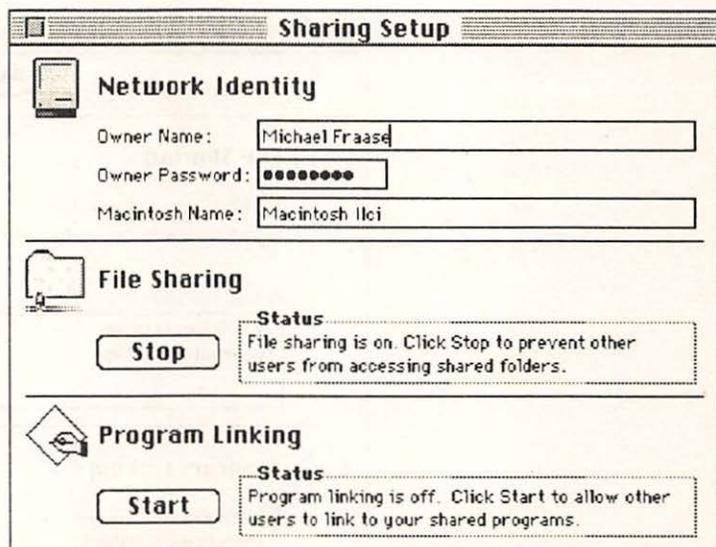


Figure 3.16 System 7.0 nondedicated file service.

Figure 3.17 shows the User Setup options that are available to you under System 7.0. It's similar in function to previous versions of AppleShare, Apple's dedicated file server. It enables you to allow other users to access the files on your system. It also enables you to allow remote users to define links to applications that reside on your hard disk.

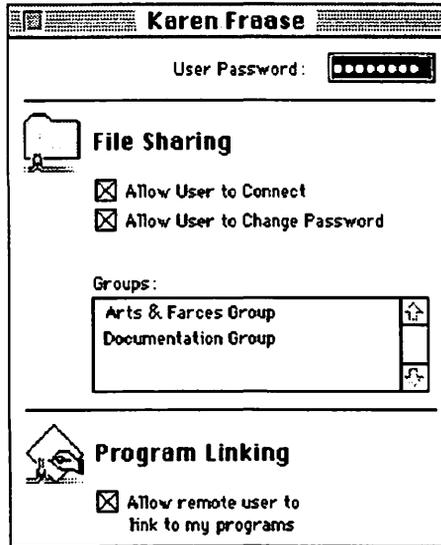


Figure 3.17 System 7.0 user privilege settings.

Distributed file server

A distributed file server, like the nondedicated file server, is a software-based network service running on each workgroup member's computer. Each workgroup member that is part of the distributed file server shares the load of file service for the network.

The main advantage of a distributed file server is that a dedicated Macintosh is not required to provide file service for the workgroup. This reduces the cost of file service for smaller workgroups. Unlike a nondedicated file server, however, a distributed file server offers the added benefit of being appropriate for larger workgroups as well.

Additionally, one of the most significant drawbacks of the nondedicated file server is addressed with a distributed server, by allowing administration to be centralized if needed.

The main drawback of a distributed file server is reduced performance in certain configurations, relative to a dedicated file server. If certain circumstances, however, a distributed file server can actually be configured to provide greater performance than a dedicated file server.

The most popular distributed file server for the Macintosh is Information Presentation Technologies Inc.'s Personal Server Network.

Electronic Mail

Electronic mail is another common software-based network service. Email allows any user to send messages and attached files to any other user on the network. Certain Email software packages allow network users to send mail messages and files to people outside the local workgroup, using telecommunications services like AppleLink, MCI Mail, and CompuServe.

When a workgroup member sends an Email message to another user, the message (and any attached documents) is stored on a centralized mail server. The recipient is notified that he or she has mail waiting and can retrieve it the next time they log into the mail server. The mail server does not require a dedicated Macintosh and some can be used in conjunction with dedicated file servers.

Electronic mail is very efficient because the recipient doesn't have to be available—their computer doesn't even have to be turned on—at the time the message is sent. Additionally, a single message can be sent to various recipients and need only be sent once.

Printer Service

A print server is a hardware- or software-based network service that lets multiple users on the network send files to the same printer at the same time. This allows the workgroup members to share expensive printers.

The printer stores the files in a queue and processes them in the order in which they were received. Some print servers allow an administrator to override the order in which the files are processed.

The main advantage of a print server is that it allows a user to return to productive work after sending the print job to the printer rather than waiting for the printer to finish its job. Network-based printer service is shown in Figure 3.18.

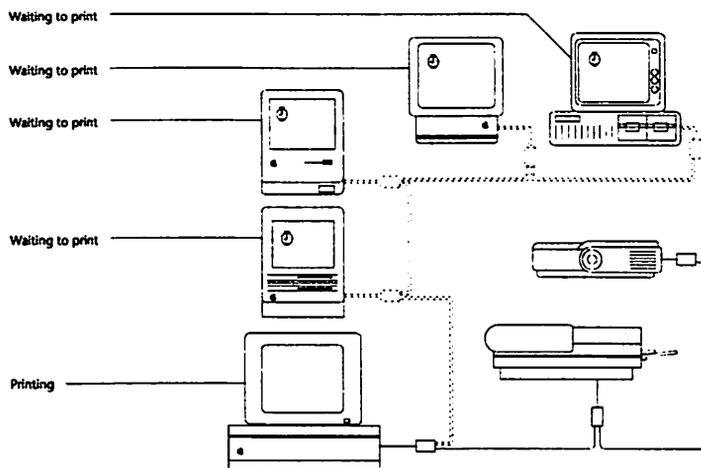


Figure 3.18 Network-based printer service.

Network Standards

A standard is a mutually agreed upon set of specifications for designing software or hardware. Standards can be set by a single vendor, a group of vendors, or a standards organization. The ISO/OSI model covered earlier is a model. It is not a standard.

Vendor-developed protocols, such as AppleTalk and IBM's Systems Network Architecture (SNA), can become de facto standards by being assimilated by the community at large and other vendors.

Industry standards, such as Ethernet and TCP/IP (Transmission Control Protocol/Internet Protocol), have usually been jointly developed by a consortium of industry organizations.

International standards, such as X.25 and X.400, are developed by international committees with representatives from industry and government.

X.25 is the CCITT interface standard for packet-switched networks.

X.400 is the CCITT international standard for electronic mail.

The Future of Networking

Future networking advancements are generally focused on providing workgroup members with a wider bandwidth. In the foreseeable future it will be possible for a workgroup to enjoy a level of bandwidth that appears to be virtually infinite. Of course appearances are deceiving and I'm sure network services will be developed to absorb every bit of available bandwidth.

I believe that the next significant step in networking will be the advent of a completely digital telephone network.

ISDN

Integrated Services Digital Network (ISDN) is a set of standards that will eventually unify voice, data, text, and image communications into a single digital wide area network. The standard is currently in flux and continues to evolve. It represents nothing less than the total conversion of the world's telephone network from analog to digital. ISDN, in simple terms, will digitize the only remaining analog portion of the telephone network: the link between your local telephone switching station and your desktop.

The ISDN standard is comprised of two components:

- *Digital network.* The digital network part of ISDN includes the ISDN lines, equipment, and software that controls the network.
- *Integrated services.* The integrated services portion of ISDN are the actual applications and services that will be developed. Such applications and services will focus on the moving of voice, data, and image information across the physical network.

There are two configurations of ISDN that will initially be available. One is geared for consumer and home use, the second is aimed at businesses.

Basic Rate Interface (BRI)

The minimal configuration is called the Basic Rate Interface (BRI) and is comprised of two digital lines (referred to as Bearer, or B, channels) and a third Delta (D) channel. This basic ISDN service is commonly referred to as "2B+D." Each Bearer, or B, channel is capable of carrying any kind of data

at a speed of 64,000 bits per second. The Delta, or D, channel is used to control the B channels and is capable of transmitting and receiving data at 16,000 bits per second.

The two "B" channels in ISDN's Basic Rate Interface result in a total available bandwidth of 144,000 bits per second. Using both "B" channels at the same time allows you, for example, to simultaneously talk with a workgroup associate while you collaborate on a document using screen-sharing software.

Primary Rate Interface (PRI)

The ISDN configuration targeted for business use is called the Primary Rate Interface (PRI). The PRI specification offers 23 B channels in North America and Japan (30 B channels in Europe) and a single D channel. All of the channels in the Primary Rate Interface version of ISDN—including the D channel—are capable of carrying any kind of data at 64,000 bits per second.

The North American/Japanese version of ISDN's Primary Rate Interface offers a total available bandwidth of 1.54 million bits per second. It's believed that the earliest users for ISDN's Primary Rate will be those businesses wanting to provide bulk connectivity into the network. Think of the ramifications for the online telecommunications services and educational institutions, to mention only two.

The ISDN "B" channels in either BRI or PRI service configuration are used for the actual user-to-user communications. Each B channel can carry any type of data and the channel may be either circuit-switched or packet-switched.

The ISDN "D" channel in either service configuration is used to carry signalling information that controls the B channels. A little-known aspect of the "D" channel is that because the signalling information does not use the complete channel at

all times, the channel can also be used for X.25 packet switching or other low-speed data transfer activities.

Initial ISDN implementations for the Macintosh will likely consist of a NuBus card that allows the Mac to physically connect to the ISDN network and specialized software to carry out communications activities.

The completely digital telephone network—ISDN—will radically change the way we communicate and collaborate. The combination of speed, bandwidth, and network of connections will make completely new kinds of software and services possible. ISDN will be the foundation for entirely new forms of collaborative applications. Individuals will be able to move text, graphics, animation, video, sound, and images to any location, at any time. Conceivably, it will no longer matter where data are stored. The implications are overwhelming. The most impressive aspect of ISDN for most users will be in the integrated services side of the equation.

There is not a general consensus about when ISDN service will be widely available. Predictions range from the mid 1990s to twenty years later. In the United States, the telephone companies have not made aggressive moves in updating the network. On the other hand, the service providers are mostly taking a wait-and-see attitude until customer demand can be demonstrated. Most customers see no need to jump on the bandwagon until they can be assured that everyone has ISDN service.

Other countries, most notably Japan, have taken a proactive stance in regard to ISDN. The Japanese are subsidizing the implementation of their ISDN system to such a degree that if you order new telephone service in Japan, you pay *more* for non-ISDN service.

Potential ISDN Integrated Services

The integrated services side of ISDN will likely fall into four general categories:

Screen-Based Telephony

Screen-based telephony integrates the telephone services we currently take for granted with the Macintosh desktop. This enhanced functionality and performance will likely result in new metaphors for communication. Look at how the Fax machine has changed some of the ways we communicate; people fax everything from nondisclosure agreements and contracts to lunch orders these days. Imagine what will happen when all of this and much more can be run from the Macintosh desktop. Most of the initial ISDN applications will use some form of screen-based telephony.

I envision the full integration of ISDN communications services with the Macintosh desktop. Instead of having to know all about baud rates, parity settings, and communications protocols to send a document to a workgroup associate, you'll be able to simply drag the document to the transfer icon on the desktop. The computer will be intelligent enough to figure out to whom it is addressed and the best method of getting it there. For instance, if it's a document that requires a signature, it will be sent as a fax. If it's a document that requires annotation it will be sent as a live, editable document. Data communications will truly become as easy as dialing the phone. Easier, in fact, because you won't even have to remember the phone number.

The picture phones featured at several World Fairs, long a gleam in the eye of the telephone companies, will become a reality. Imagine the boon in the "900" services.

Imaging

Imaging applications will be combined with data-compression technology and ISDN to provide huge networks that connect newspapers, magazines, book publishers, and other print media to service bureaus, illustrators, photographers, designers, editors, and writers.

The average color news photograph is extremely large—multi-gigabyte images are not uncommon—and the publication process sometimes requires up to four or five transitions from analog to digital information. New imaging technologies combined with ISDN links will cut the transitions to two points: the sender (the photographer) and the recipient (the publication). More efficient compression technologies will do virtually the same thing for video and radio broadcasting.

The dark side to this process is the whole can of works currently being explored surrounding digital retouching, and similar technologies. A discussion of the societal repercussions of this aspect of the technology is beyond the scope of this work, but it is important and must be attended to sooner rather than later.

Resource Sharing

One of the central benefits of local and wide area networks is the possibility of sharing resources. ISDN technology holds tremendous promise in this area.

Current resource sharing applications for the Macintosh include Farallon Computing's Timbuktu and Timbuktu/Remote and Microcom's Carbon Copy Mac. These are screen sharing applications (covered in-depth in Chapters 4 and 5) that allow two or more workgroup members to work collaboratively by sharing a common screen space.

ISDN holds the potential of enhancing applications like these to include the sharing of animation, sound, video, and other desktop resources in addition to the screen sharing we currently enjoy. This will allow workgroup members to share just about anything they can create with (or bring into) their personal computers.

The logical extension of resource sharing is full-motion video conferencing, already being demonstrated by Apple and third-party developers.

Things We Can't Even Imagine Yet

In the relatively near-term future, we will be able to employ an intelligent software "agent" to search through remote databases, archives, video libraries, etc. for information that is relevant to our work or personal lives. This will allow the long-touted personal newspaper to become a reality and the completeness offered by access to huge amounts of information will be overwhelming without the help of agenting software.

Virtual reality will enjoy a big boost from ISDN technology. No longer restricted to the virtual worlds inside our own computers, we will have access to a vast network of other virtual worlds and will be able to collaborate within them. They'll also be incredible stress relievers.

Apple Computer identifies one of the most significant applications for ISDN technology as object communication. Object communication will allow computers to easily call up other computers and exchange software objects. This will allow a highly evolved form of interprocess communication, called process-to-process communication by Apple.

Fiber Optic

Some observers believe that we should bypass ISDN altogether and replace the entire telephone network with fiber optic cable. The argument is sound if for no other reason than it will eventually have to be done anyway. ISDN, according to proponents of fiber optic technology, does nothing more than squeeze every bit of performance out of a quickly aging copper network.

Fiber optic cable carries data at the speed of light. Where ISDN is capable of transmitting data at speeds of up to 14 Kilobits per second, fiber optic will move the same data at a speed of 20 Gigabits per second. That's 20 billion bits of information per second if you're keeping score.

The problem with fiber optic is the same chicken-and-egg situation that confronts ISDN, only worse because it's marginally more expensive. The telephone companies don't want to bear the burden of investing in fiber until there are applications. The developers don't want to invest in fiber optic application development until the network is in place.

Some believe that the telephone companies should be allowed to act as information providers—something that has been forbidden since the Great Telecommunications Dives-titure—in order for fiber optic technology to mature into a viable medium. Opponents argue that such a move would give the telephone companies a monopoly on the content carried across the entire fiber optic bandwidth.

In any case, for fiber optic to be successful, companies that are heavily invested in older technologies—specifically the cable television and traditional publishing industries—will have to come around to seeing how fiber optic can be beneficial to them. Right now these businesses aren't looking forward to drastic changes in the way they do business.

Because of fiber optic's incredible bandwidth, current legislation dictates that the installer of the network (presumably

the telephone company) has no power over who has access to the fibers contained within the cable. All information providers, large and small, would have to compete in an open marketplace.

Pacific Bell has been investigating fiber optic options in its technology lab. The company has a mock fiber optic telephone network installed and is able to run multiple applications simultaneously in real-time. Some of the applications, like video teleconferencing, are fairly predictable. Other applications, however, point to the capability of the medium. LiberNet, for example is converting traditional paper books to digital volumes, accessible over the fiber optic telephone line. Another application, Switched Multimegabit Data Service, offers a public-access local area network between business sites. The service eliminates the need for dedicated data lines and lets businesses pay for only the connect-time actually used.

Pacific Bell and other fiber optic researchers are attempting to show that anything ISDN can do, fiber optic can do better. So far, all indications are that they may be right.

Managing the Network

Managing a network, even a small one, is a task best approached with patience and diligence. It's not an undaunting task, but it's a manageable one if approached as a series of procedures. Current needs must be assessed, and plans laid for managing growth; few networks stay the same and can be best approached of as a series of evolutions.

Assessing your Networking Needs

Assessing your networking needs sounds complicated and time consuming. It doesn't have to be complicated, but you

are going to have to invest some time in the process. If you spend a reasonable amount of time assessing your needs now, you won't be confronted with dead-ends later on. The assessment process can be broken down into a series of steps that make the process manageable.

Inventory Your Current Needs

Begin by taking an inventory of what you perceive to be your current network needs. You can approach this in an open, freewheeling manner or you can perform a thorough needs analysis.

In either case, the number of workgroup members and the number of devices that need to be connected are key factors in this inventory. Also, be sure to take into consideration any workgroup member or members that are located away from the main site.

At this time it's too soon to begin assessing the various network media available. Instead focus on who and what will need to be connected; worry about the how later.

Determine Network Service Needs

Next, determine your network service needs. In the previous section, various network services were discussed briefly. This is the time to begin to determine which ones are suitable for your network.

Begin this determination process by asking yourself and your coworkers what types of communication tasks the network will be used for. Will you need file service in addition to file sharing, or will an electronic mail system that allows documents to be attached suffice? Will you need a print server? What other sorts of network services does your work-

group anticipate? This is the time to get a handle on these questions.

Evaluate Solutions

The next step is to evaluate the various solutions that are available and appropriate for your workgroup situation. Important considerations at this stage include an evaluation of various product offerings and vendors. It's important to include an intelligent evaluation of potential vendors. Questions to ask yourself are, How well does the vendor adhere to established standards? How compatible are their products with those from other vendors? Will the vendor be in business several years into the future.

Product evaluations are provided for you in the following sections continuing through Chapter 4: Groupware Implementations. Tutorials for actually implementing various network services are provided in Chapter 5: Tutorials. Use these sections as a starting point in your assessment process. This is the time to determine which network services you want to provide for your workgroup and which media best supports those services.

Anticipate and Prepare for Growth

After evaluating the solutions that are available, anticipate and prepare for growth of your workgroup's network. If you spend a little bit of time and effort preparing for the inevitable growth of your network, you will be playing a proactive role in a smooth network evolution rather than constantly trying to play a reactive, catch-up role.

Begin by realistically projecting the needed expansion for your network. Will you be adding workgroup members in the foreseeable future? What about moving from a nonded-

icated file server to a dedicated one? Will your choice of an electronic mail system be flexible enough to grow with your needs? Do you anticipate greater network traffic in the relatively near term? Answer these questions the best you can now; you'll reap the dividends later.

Perhaps the best approach is to keep your network as flexible as possible. The computer industry has always experienced drastic changes, and there's little indication that this will not be the case in the foreseeable future. Planning flexible networks that can adapt to multiple protocols will be more likely to keep pace with evolving technologies.

Decide on the Physical Layout

Finally, decide on the physical layout of the network. The topology of your network and its physical layout will determine how easily you will be able to add new network services, devices, and users in the future.

The topology you choose will also determine how easily you will be able to connect additional networks to form an internet.

Refer back to the Networking Basics section of this chapter for information on the various network topologies to determine which one is the most appropriate for your workgroup.

Administering the Network

There are two basic approaches to network administration. You can either appoint a single person to be responsible for the smooth operation and updating of the network or you can distribute the tasks to various members of the workgroup.

My investigations and experience have led me to believe that a sort of hybrid of the two extremes tends to work best.

Appoint a single person as a network administrator to oversee and maintain the network, especially if you're using a dedicated file server. Instruct individual workgroup members in how to maintain their own workspaces, especially if you're using a distributed file server. If you're using a dedicated server, workgroup members should be instructed in how to maintain the shared file space. In larger workgroups, it may be necessary to organize an administration group to manage the network.

As an example, consider the scenario of routine software upgrades with an enterprise-wide network of 100 or more users. Imagine the time commitment involved with updating hundreds of separate workstations, especially if the software requires a specific serial number installation for each workstation. A simple task that should require no more than 15 minutes suddenly expands to fill days of someone's time. Tracking down original disks, sorting serial numbers, and matching the update disks to the correct workstation grows more complicated as the installed base of networked users grows.

A network administrator, then, is responsible for maintaining the day-to-day functionality of the network. During installation phases, the administrator's job will be to install or move cabling, connecting and configuring devices, and installing and configuring software.

After the network is installed, configured, and up and running, the administrator's job is to manage the network services and troubleshoot any problems that may occur. Specialized software is available for these tasks and will be covered in-depth in Chapter 4: Groupware Implementations. Briefly, this software can automatically update user software (avoiding the scenario mentioned above) as well as report on the number, location, and status of various network devices and services. Diagnostic software is available

that can identify a variety of problem conditions on the network and suggest solutions. Other software is available to monitor network performance and analyze the demands and level of use throughout the network.

The final aspect of network administration is day-to-day maintenance of the network and its services. If a user accidentally disconnects his or her LocalTalk cable, for example, the job of discovering and correcting the situation usually falls on the network administrator.

The goal of effective network administration is to provide the workgroup members with as transparent a network as possible with the highest level of performance and lowest level of disruption possible.

Collaborative Computing in the Mixed Platform Network

In the near-term future, the hardware platform decision will be easier to make, if only because it will be less important. Applications that are designed for multiple hardware platforms are becoming widely available, and this is a trend that will continue for the foreseeable future. The concept of document import and export filters makes it possible for disparate applications on a variety of computing platforms to share data files with each other. What will become more important for workgroups of all sizes is integrating various hardware components into the network.

Setting up the Mixed Platform Network

For the sake of brevity the remainder of this book (except where noted) assumes that your workgroup and its needs are fairly typical. The typical workgroup consists of five to seven members with the potential for interconnection with other workgroups within the enterprise in the future. The follow-

ing examples will assume a workgroup comprised of five people with an expectation of adding two more within the next 12 months.

The remainder of this book assumes that your workgroup needs are fairly typical. You'll need file service, printer service, electronic mail service, and screen-sharing capabilities. It would be nice to provide room for ancillary networking products, such as meeting scheduling, when they become available. Each of the five members within your workgroup has been using a Macintosh for at least three months. Your workgroup's computing resources consist of the following:

- 1 Macintosh SE/30
- 2 Macintosh IIsx
- 1 Macintosh IIfx
- 1 IBM PS/2
- 1 LaserWriter IINT

Additionally, your budget includes resources for adding two more workgroup members within the next 12 months, one of whom will be using a Macintosh IIfx; the other a NeXT-cube.

Your workgroup, then, will eventually included three disparate computing environments: Macintosh, IBM PC, and NeXT.

Two Problems for Mixed Platform Networks

Two significant problems for mixed platform networks exist that must be addressed prior to installation of the network and software. These are the exchange and translation of files and network performance degradation caused by large graphic files. Each of these problems will be addressed separately in the following sections.

Sending and Translating Files

The simplest way to transfer text and graphics between computing platforms is to choose software that exists for all platforms.

For Mac-to-PC Connections:

- Microsoft Word for word processing
- Microsoft Excel for spreadsheet work
- Microsoft PowerPoint for presentations
- Microsoft Works for an integrated application
- Aldus PageMaker for page layout
- Aldus Persuasion for presentations
- Adobe Illustrator for PostScript graphics
- QuarkXPress for page layout
- Ventura Publisher for page layout

For Mac-to-NeXT Connections:

- T/Maker WriteNow for word processing
- Frame Technology FrameMaker for page layout
- Informix WingZ for spreadsheet use
- QuarkXPress for page layout
- WordPerfect for word processing

Text compatibility is best maintained (if you can't use the same application on all platforms) by saving files in the rich text format (RTF).

Bit-mapped graphics should generally be saved as Tagged Interchange Format files (TIFF).

Object-oriented graphics should be saved as Encapsulated PostScript (EPS) files.

Font compatibility is a primary issue. You've got to have the same fonts installed on each platform. Currently, there is no translation to or from TrueType, LaserJet III, or Adobe PostScript Type 1. If the PCs are used only for text, however, the font compatibility issue is not as severe.

Large Graphic Files Degrade Network Performance

A LocalTalk network is an excellent solution provided that network traffic is relatively low enough. If large graphics files must be moved across the network, however, be aware that the network will bog down quickly. In these cases, the best current alternative is to install high performance 32-bit Ethernet cards on all machines that will handle large files.

In practice, the best solution is to keep extremely large files off the network as much as possible. Many design studios and service bureaus use a single image manipulation workstation connected directly to the imagesetter or other output device.

Macintosh-to-PC

The most appropriate workgroup solution for Macintosh-to-IBM PC networking is to connect both machines to a LocalTalk network. Installing a LocalTalk board in the PC and connecting to AppleShare or AppleShare-compatible file servers such as Personal Server Network is the simplest solution for most workgroups. This arrangement also allows the PCs access to any Macintosh laser printers available on the network. A typical workgroup setting that integrates IBM

PCs or compatibles within the LocalTalk network is shown in Figure 3.19.

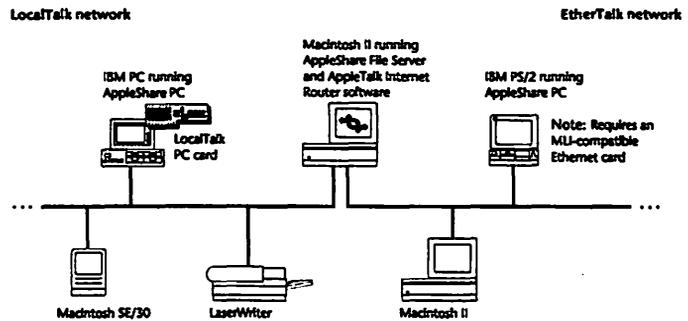


Figure 3.19 LocalTalk/EtherTalk network with IBM PC integration.

Figure 3.19 shows an expanded workgroup configuration comprised of both LocalTalk and EtherTalk networks. The PC attached to the Ethernet network requires an Ethernet card and an appropriate router—either hardware or software—is also required.

If network traffic warrants, the entire local area network can be easily converted to Ethernet by installing appropriate hardware cards in each of the Macintoshes on the network. Such a configuration would not require a router.

Some workgroups may find it unnecessary to connect the PC to the local area network. This may be a result of the PC user not needing file or printer services across the network or environments where the PC user or users are remotely located. For local PC users that only require occasional file transfers an adequate alternative may be to use DataViz's MacLink Plus/PC. For remote PC users, dial in access to an electronic mail, conferencing, or bulletin board system may be a more appropriate solution.

Macintosh-to-NeXT

Macintosh-to-NeXT connectivity is more complicated because at the time of this writing, no LocalTalk cards for any of the NeXT models are available. Several options are in the development stage, but a better alternative is to use straight Ethernet (since Ethernet is built into every NeXT computer) or bridge a LocalTalk network for the Macintoshes with an Ethernet network for the NeXTs.

Many mixed-platform networks will require transparent file sharing between platforms. The best alternative available for this situation is to use a Cayman Systems GatorBox. This hardware device provides a LocalTalk to TCP/IP Ethernet gateway between the two platforms. Cayman also offers specialized software, GatorShare, to provide transparent AppleTalk File Protocol (AFP) to Network File System (NFS) file sharing.

Figure 3.20 shows an expanded mixed-platform network configuration comprised of Macintoshes, an IBM PC, and a NeXTcube. The Macintoshes and PC run on LocalTalk while the NeXTcube runs on Ethernet and the two networks are bridged with a GatorBox.

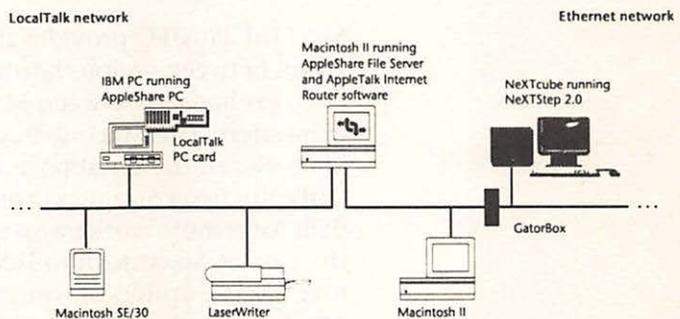


Figure 3.20 LocalTalk/Ethernet network with IBM PC and NeXT integration.

At a total price of more than \$6,000 (\$3700 for the Gator-Box; \$2400 for the GatorShare software) this is an expensive solution for small workgroups. An additional drawback—quite significant in certain environments—is that the Macintosh can see NFS volumes, but the NFS host cannot see Macintosh volumes.

Network Alternatives

In the smallest workgroup situations, a network may not be the best solution. In small workgroups where collaborators don't use the same applications across all platforms, a cable can be connected between the serial ports of dissimilar computers. Data translation software, such as MacLink Plus/PC can be used. Additionally, MacLink Plus/PC contains a set of translators (MacLink Plus/Translators) that can be used with some programs to translate files transparently.

DataViz Inc.'s MacLink Plus/PC is an appropriate for small workgroups that do not need file sharing or file service among multiple platforms. If, as in the case of the sample workgroup specified here, your workgroup is comprised mostly of Macintoshes with one or two IBM PCs or NeXT workstations, MacLink Plus may be all you need.

MacLink Plus/PC provides file transfer and automatic translation between various hardware platforms and enables files to be exchanged between Macintosh, IBM PC, NeXT, or Sun computers. The Macintosh controls a connection with a cable between the serial ports of two disparate computers. Only the two computers connected via the cable (or a modem for remote workgroup members) can exchange files. In the case of Macintosh to IBM PC connections, the IBM PC user has the option of controlling the file exchange session. All of the foreign computers supported by MacLink Plus/PC work in the same manner, and connection files are provided for various hardware platforms. This section will use a Macintosh to NeXTcube connection as an example.

A specialized cable is included with the MacLink Plus/PC distribution package and allows a Macintosh to connect to virtually any modem or other computer's serial port with a set of cable adapters. Most small workgroups will find it useful to leave the cable connected between the two computers and use a single Macintosh connected to the NeXTcube for all file transfers between the two platforms.

The file transfer and translation process is straightforward. A specialized application, DataVizBridge, is included with every NeXT computer sold. Launch the DataVizBridge application on the NeXT and click the Connect button shown in Figure 3.21.

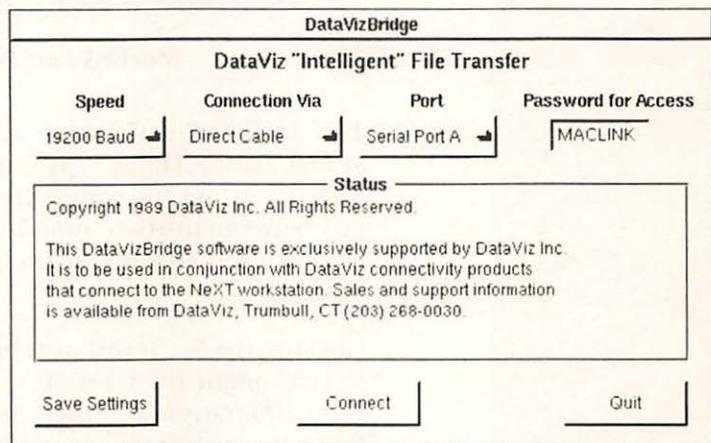


Figure 3.21 *DataVizBridge for the NeXT.*

Next, launch the MacLink Plus/PC application on the Macintosh connected to the NeXTcube. Open the document named "Link to NeXT." You'll see a display similar to the one shown in Figure 3.22.

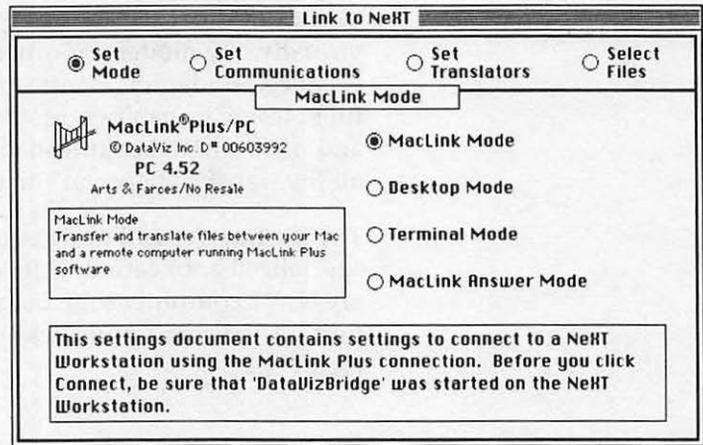


Figure 3.22 MacLink Plus/PC "Link to NeXT" main window.

The MacLink Plus/PC main window looks quite complicated at first glance. Don't let it intimidate you. The "Link to NeXT" settings file comes pre-configured for the serial port link between the two computers. The only options you need to be concerned with are the Set Translators and Select Files buttons.

Clicking the Set Translators button allows you to configure the automatic file translation feature of MacLink Plus/PC. Over 100 translation filters are available and include translators for most word processing, database, spreadsheet, and graphics applications.

The appropriate setting for translating Encapsulated PostScript (EPS) images from the NeXT to the Macintosh is shown in Figure 3.23.

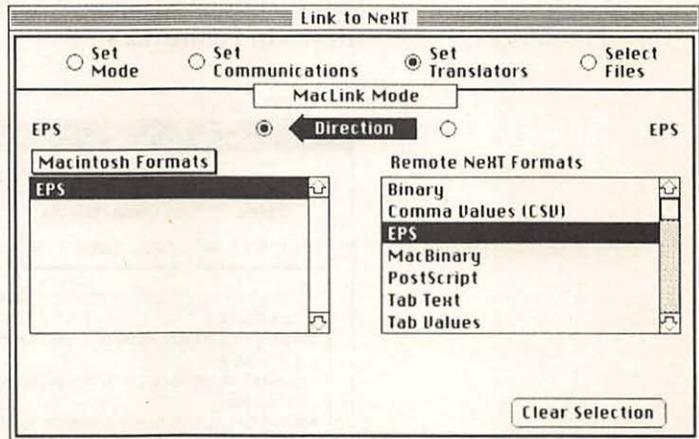


Figure 3.23 *MacLink Plus/PC Set Translators settings for EPS images.*

Since all activity is controlled from the Macintosh side of the connection, note that a Direction button is provided to allow you to transfer files in either direction.

The next step is to select the files to be transferred. You do this by clicking the Select Files button shown in Figure 3.23. Clicking the Connect button establishes the actual connection with the NeXT. The connection is immediate and the display changes to a listing of the local Macintosh file hierarchy in the left panel and the remote NeXT file hierarchy in the right panel of the window.

NeXT directories and subdirectories are displayed as folders; to navigate you simply double-click to open a folder in a manner identical to the standard Macintosh file dialog. The actual transfer is initiated by clicking the Transfer button.

MacLink Plus/PC on the Macintosh side and DataVizBridge on the NeXT side each maintain a log of all transfer and translation activities. The Macintosh side writes the log to a

text file on disk, while the NeXT side displays the log in the DataVizBridge window. A sample DataVizBridge log is shown in Figure 3.24.

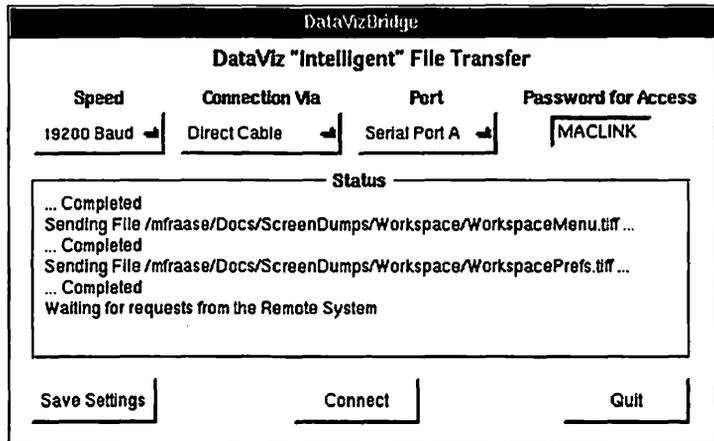


Figure 3.24 DataVizBridge file transfer and translation log.

Because the NeXT is capable of background processing, most small workgroups will find it beneficial to leave DataViz-Bridge running in the background.

Chapter Four

Groupware Implementations

Robert Johansen, in his *Groupware: Computer Support for Business Teams*, identifies 17 approaches to groupware:

- *Face-to-Face Meeting Facilitation.* Face-to-face meeting facilitation is the computer-supported equivalent of flip-chart pads and felt-tip pens. Possible extensions include extensive technological support for traditional meeting facilitation.
- *Group Decision Support.* Peter Keen and Michael Morton, in *Decision Support Systems*, identified decision support systems using computers to “(1) assist managers in their decision processes in semi-structured tasks; (2) support, rather than replace, managerial judgement; and (3) improve the effectiveness of decision making, rather than its efficiency.” Johansen advocates extending the decision support metaphor into group decision support systems. Systems such as these have been in use since the 1970s in various forms.
- *Computer-Based Telephony Extensions.* The telephone has the tremendous advantage for collaboration of being familiar. Everyone understands a telephone and knows how to operate it. Johansen sees two basic approaches to telephony that can overlap and push collaborative

computing to the forefront. The telephone network itself and the private branch exchange (PBX) can both be piggy-backed to accommodate collaborative computing.

- Northern Telecom's Meridian system provides a service called Meeting Communications Services. This service allows the equivalent of a conference call with visual aids displayed on each participant's display screen.
- Similarly, AT&T's Alliance service provides conference calls with the ability to exchange graphics between participants.
- *Presentation Support.* Computers can be used effectively to create and present slide shows. Macintosh software such as Microsoft's PowerPoint, Aldus' Persuasion, and Symantec's MORE are marketed specifically for this purpose and are quite effective when paired with projection screen technologies.
- *Project Management.* Workgroups have obvious needs for project management software, and such tools are well established in the personal computer market. Because workgroup members have more important things to do than project management, this software must be easy to use by everyone involved.
- *Calendar Management.* Workgroup members must coordinate schedules on a frequent and ongoing basis. Ideally, this software would work in a manner that would allow each workgroup member to designate time that he or she is available with a series of weight-values. The software would establish a common block of available time and would resort to the weight-values in the event that such a time block were unavailable. Unfortunately, most group calendar management software is currently implemented in such a way that the needs of all workgroup members are not taken into account and weight-values are assigned in a pecking-order fashion.

- *Group-Authoring.* Group-authoring tools are likely to be the first breakthrough groupware products for the simple reason that most of us already author documents in a group setting and that behavior is easily translated into the electronic medium. Most documents get passed around to colleagues for comment prior to publication and distribution. The copy is received back with comments penciled in the margin. The author then integrates these comments into the document as needed. Computer-based group-authoring tools must be implemented in a way that allows the system to keep track of all revisions, allowing workgroup members to suggest changes without altering the original document.
- *Computer-Supported Face-to-Face Meetings.* Computer-supported face-to-face meetings require the electronic equivalent of the familiar white board that lines a wall in most meeting rooms. Such a setup requires the presence of a workstation for each group member and a shared space that is displayed outside of each member's personal space. The software must be capable of providing direct group support and must be simple enough to use without extensive training. Mark Stefik's work at Xerox's Palo Alto Research Center (PARC), specifically the PARC COLAB, is a fundamental implementation of this groupware concept.
- *Screen-Sharing.* Anything that can be displayed on an individual workgroup member's screen should be able to be shared with other members of the workgroup. Xerox PARC researchers have labelled this type of software WYSIWIS (what you see is what I see). Macintosh-specific implementations of this concept include Farallon Computing's Timbuktu and Timbuktu/Remote and Microcom's Carbon Copy Mac.
- *Computer Conferencing.* Electronic mail systems are designed for individual-to-individual messaging. A computer conferencing system would allow groups to

access a shared message base. Various commercial conferencing systems have been implemented.

- **Text Filtering.** Text filtering software allows workgroup members to sift through vast amounts of free-form or semi-structured text. More powerful filtering can be accomplished by applying greater structure to the search criteria and filtering mechanism. Common Macintosh implementations of text filtering include Microlytics' GoFer and On Technology's On Location.
- **On Location.** GoFer and On Location are woefully inadequate for what Johansen predicts for the future generation of text filtering tools. "Text filtering can also be used to identify people with common interests. In this way, text filtering can be used for computer support of much larger communities, creating a kind of magnet for filtering text. The team uses the filter to search out information and people that can help move its task forward. As is typical with many work teams, the members are working in a field that is still not mapped or well understood; they are ahead of the key words in traditional databases. The filter helps them specify just what kinds of information they want. Each morning, the filter prints a personalized 'newspaper' for each team member, showing items from the preceding day's news, as well as new findings from the ongoing search for leads. Person-to-person messages are also filtered to insulate the team members from low-priority interruptions."
- **Computer-Supported Audio/Video Teleconferencing.** Apple has demonstrated a video conferencing system that operates on existing Macintosh platforms and employs a broadband local area network capable of sustaining video information.
- **Conversational Structuring.** Conversational structuring in a collaborative computing environment is best ex-

plained as a series of prompts and responses between the workgroup member and his or her personal computer. Johansen uses the following conversational structuring example to illustrate the point.

Workgroup member: "OK, let's do it."

Computer prompt: "Who should do what?"

Workgroup member: "I guess I should get it going."

Computer prompt: "What will you do, when?"

Workgroup member: "I'll do it by Friday."

Computer prompt: "I'll put it on your calendar and advise the rest of the team. What, exactly, are you agreeing to do?"

- *Group Memory Management.* All workgroups have clear need for a group memory system. Each workgroup member must be afforded access to the group memory and tools should be available allowing the memory to be searched quickly and intuitively. The hypertext model is an excellent group memory management system because of its inherent ability for nonlinear access to information and associative trails through immense bodies of information.
- *Spontaneous Interaction.* Spontaneous interaction must be supported in any collaborative environment. In the collaborative computing environment, this is the electronic equivalent of hallways and the water cooler.
- The best example of a spontaneous interaction model in the collaborative computing environment is Xerox PARC's System Concepts Laboratory (SCL). SCL was comprised of two halves, one in Portland, Ore. and the other in Palo Alto, Calif. Audio, video, and data links were maintained between the two sites and were available to all workgroup members on a 24-hour-a-day basis. The intention was to create an electronic environment that would encourage spontaneous interaction.

- *Comprehensive Workgroup Support.* Support services will form a central part of the groupware process. Ask a network manager of a major Macintosh installation about the nightmares of perpetual software upgrades. Imagine having to install several hundred serialized software upgrades. A process that should take 15 minutes, if automated, currently takes days.
 - Other forms of comprehensive support will also be required. According to Johansen, the "general direction is toward putting users 'inside' their computing environments." Most systems like the one envisioned by Johansen have their foundation in the early work of Douglas Engelbart. Engelbart built a prototype system in the early 1960s called NLS (oNLine System). A commercial version of NLS has been available from McDonnell Douglas but has been relatively unsuccessful because it remains somewhat cumbersome to use. Mountain View, Calif.-based Metaphor Computer Systems has developed an integrated computer system geared specifically to providing comprehensive workgroup support.
- *Non-Human Meeting Participants.* In the relatively near-term future, computer software should be able to function, at least in a limited fashion, as members of a collaborative workgroup. Most observers recognize this as the domain of artificial intelligence, but agents are much more likely to be realized in a more realistic time frame. There are no current examples, although again, CompuServe Navigator comes close to being an unintelligent implementation.

Using the Right Tool for the Job

Good tools are easily identified. They tend to feel like extensions of yourself that help you accomplish tasks in a more

efficient manner than you would otherwise be capable of. Using the proper tools are essential to performing any task, yet we ask many of our knowledge workers to use inferior tools, or worse, none at all. Tools are what we use to add value in our work.

According to Michael Schrage, "Tools are, literally and figuratively, the way people come to grips with their work. Too often, organizations apply a perverse calculus of counterproductivity by asserting that technology equals machines equals tools. These organizations try to substitute technology for blue-collar labor and white-collar thought. Or, if they consider themselves enlightened, they position new technology as an amplifier, a complementary prosthetic for boosting managerial and work-force productivity."

The right tool is what Ivan Illich has identified as a "convivial tool." Ivan Illich: "Tools foster conviviality to the extent to which they can be easily used, by anybody, as often or as seldom as desired for the accomplishment of a purpose chosen by the user. The use of such tools by one person does not restrain another from using them equally. They do not require previous certification of the user. Their existence does not impose any obligation to use them. They allow the user to express his meaning in *action*."

When the Right Tool Isn't Always the Right Tool

The Right Tool for individual work is not necessarily the Right Tool for collaborative work. The most common office tools that we use on a daily basis were designed for individual use, not collaborative use.

The very availability of a certain set of tools can cause people to behave in certain ways. A laser printer or imagesetter in your office encourages a form of communication that a trip to a service bureau doesn't. Likewise, tools that are designed specifically for collaborative use—or at least with collabora-

tion in mind—will tend to enhance the collaborative process itself.

Language is the basic tool for the collaborative process, and the model upon which other convivial tools for collaboration should be built. This is self-evident in the tools we make daily use of. The telephone, word processor, electronic mail, answering machine, photocopier, optical character recognition, etc. are all based on language.

Michael Schrage has identified at least four ways that language can shape thought:

- *The words themselves.* If a word we need to express ourselves doesn't exist, we simply invent it. Look at how the computer industry alone has changed our vocabulary in a very short time.
- *The demanded inclusion of certain information.* Schrage uses the illustration of verb tense in Western, Indo-Germanic languages forcing us to be conscious of the passing of time. He goes on to state that the Japanese language requires the inclusion of the relative rank in the societal pecking order in every sentence.
- *Patterns of expression.* Schrage provides two wonderful examples: "English's subject-verb-object sentence structure encourages actor-action-receiver thought patterns and cause-and-effect expression. In Navaho, it is easiest to name something by describing its characteristic behavior: a duck is *naal'eethi*, "that which floats all around." The majority of Navaho nouns are made from verbs in this pattern, as a result, Navaho tends to be much richer in active descriptions than English."
- *Relative difficulty of expression.* Our language shapes thought by making certain things harder to say than others. Schrage points out that in English we generally don't provide the source of our information, where the Hopi language manages "attribution with the same facility that English speakers handle time."

Shared Space: A New Model for Collaboration

Our current models for collaboration are constrained, but useful nonetheless. The constraints are a result of following a conversational model rather than a collaborative model. The blackboard is still our most useful collaborative tool, but not for long. The blackboard has remained unchanged for six centuries and remains useful because of its high level of reliability and resilience.

Michael Schrage emphasizes the importance of augmenting our existing tools to arrive at appropriate models for collaboration. *“What would it mean if the power and versatility of these simple tools for collaboration could be amplified ten-fold? A hundred-fold? A thousand-fold? What if technology could augment the process of collaboration with the ease that a pocket calculator augments computation? What new kinds of conversation and collaboration would occur? How would conversation and collaboration be different? What new insights into creativity and discovery would these new tools yield?”*

Schrage goes on to point out the problems with models for collaborative tools by pointing out the dilemmas of managing the traffic of conversations, “the strictures of taking turns, interrupting, and maintaining conversational flow.” Conversations are invisible and fleeting; it’s hard to keep track of who’s saying what within a group. In addition, “conversations don’t have memories; only their participants do.” The serial nature of conversation—the taking turns, interruptions, and flow management—is a barricade to the act of collaboration. “In most conversations,” according to Schrage, “people take turns exchanging information, not sharing it.”

The concept of a shared space is vital to the success of collaborative tools. In turn, the shared space shapes the collaborative process. Workgroup members are free to communicate directly with each other or through the

shared space without being bothered with managing the conversation flow. This allows ideas to be shared rather than exchanged and leads to shared understandings.

According to Schrage, "Shared space heals the rift between spoken language and visual language. In our culture, we've divorced representation from human interaction. People treat speech and writing—or speech and image—as binary, either/or, competitive with each other."

Xerox Colab

Although the best shared space most of us have access to is the network of personal computers, one of the best large-scale implementations of shared space is Xerox's Colab. Michael Schrage quotes Mark Stefik, a research fellow at Xerox PARC and director of the Colab: "Collaborative computing will be much, much more pervasive than personal computing because while not everyone needs a personal computer, virtually everyone needs to collaborate."

Colab is designed as a meeting room with a group of conference tables arranged together to form a semicircle. Each workgroup member has access to his or her own personal computer. The semicircular conference table arrangement faces a large screen that displays computer data and forms the tangible (or at least visible) portion of the shared space. Workgroup members are able to send information from their computer directly to another workgroup member or to the large screen for use by everyone else.

The Colab's large screen can be divided into multiple windows and each window can be stored, resized, discarded, moved, and linked to other windows. A facility is also provided that allows workgroup members to point at windows and objects on the shared screen.

Michael Schrage points out the central benefits of shared space in general and Colab in particular: "This technology

completely changes the contexts of interaction. For one thing, a conventional conversation normally has rules of etiquette that govern turn taking. These rules evolved around the constraints of an oral meeting, where only one person can speak at a time lest the conversation degenerate into babble. But in the environs of Colab and shared space, there are visual channels that can either augment or conflict with the spoken word. Conversation isn't the only activity going on; it's not the only domain of interaction. In ordinary conversation, a speaker responds in some fashion to the previous speaker's comments; in this new environment, people may feel more compelled than at 'ordinary' meetings to respond to something that appears on the screen. Traditional notions of conversational etiquette go out the window (pun intended) if one person writes a controversial message on the community screen while another talks about something else."

Xerox Cognoter

Cognoter is a meeting tool that evolved out of the Colab experience at Xerox PARC. Cognoter enables a workgroup to create a shared outline based on the group's consensus that is useful for organizing ideas. A Cognoter session results in an annotated outline that reflects the consensus reached in the collaborative session.

The Cognoter Colab session is structured into four distinct phases or parts.

- *Brainstorming.* A Colab brainstorming session is, by all accounts, the opposite of a standard meeting. Instead of patiently waiting your turn, all you have to do in the Colab brainstorming session is to stake out a piece of the shared space and type your ideas. Since workgroup members can see all the contributions, conversations tend to evolve and items get annotated and expanded.

- *Organizing.* After the ideas are fleshed out in the brainstorming phase, workgroup members enter the organizing phase where ideas are categorized, linked, and cross-referenced to items in someone else's window or portion of the shared space.
- *Evaluating.* The ideas are evaluated in this phase of the Colab session and relationships between ideas are represented by workgroup members actually drawing arrows between related ideas within the shared space. It becomes obvious to each workgroup member which ideas are inappropriate as the ideas are ranked in order of importance, Unimportant, overly tangential, or inappropriate ideas have no links and can easily be removed through consensus.
- *Generating Output.* The final outline is printed on a laser printer and each workgroup member leaves the session with a tangible record of the collaboration.

Xerox Argnoter

Argnoter is another Colab software tool that is best described as an "argumentation spreadsheet." Michael Schrage points out that Argnoter evolved out of a recognition on the part of Xerox PARC researchers that most "misunderstandings and disputes derive from three main sources: personal positions, unstated assumptions, and unstated criteria." Argnoter makes all of these problem sources visible by placing them in the shared space.

Like Cognoter, the Argnoter session is also divided into distinct phases or parts. Where Cognoter employs four phases, Argnoter uses three.

- Proposals
- Arguments
- Evaluation

The three phases of the Argnoter session are roughly equivalent to the first three phases of the Cognoter session.

The most significant facet of both the Cognoter and Argnoter software tools within the Colab environment is the idea of parallel conversations and the underlying fact that interruption—at least as we generally recognize it—cannot occur. Everyone is free to talk at once while being assured of also being heard.

Perhaps the most glaring problem with the Colab environment is the tensions that arise over turf wars for screen real estate. Physical space and relative size tend to be seen as a measure of importance. Battles sometimes rage for prime screen areas because larger windows are seen as being more important than small windows.

Shared space, Colab, Cognoter, and Argnoter are all easily emulated within the Macintosh environment by piecing together existing applications to form a synergistic space that was certainly not anticipated by the creators of the Macintosh.

Farallon's Timbuktu, for example, can easily be pressed into duty as the backbone for a Macintosh-configured shared space. Video display cards from several manufacturers are capable of driving large-screen monitors that will form the large, white-board-like, space described above.

An outliner or idea processor such as Symantec's MORE can be used as to emulate a Cognoter session when used in combination with Timbuktu. Each member can maintain a separate window within the shared space as shown in Figure 4.1.

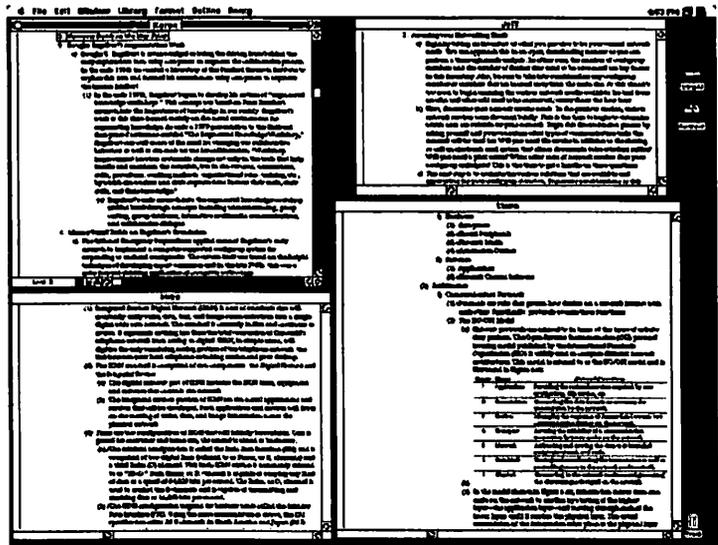


Figure 4.1 Using Farallon Computing's Timbuktu and Sy-mantec's MORE to emulate a four work-group member Cognoter session.

The Cognoter mockup implementation shown in Figure 4.1 is slow, but it works effectively. As the fundamental ideas of groupware and collaborative computing continue to evolve, more products will become available that will make the seams disappear.

Establishing and Connecting the Workgroup

For the sake of brevity the following sections assume that your workgroup and its needs are fairly typical. The typical workgroup consists of five to seven members with the potential for interconnection with other workgroups within the enterprise in the future. There are five people within

your workgroup and you expect to add two more within the next 12 months.

Your workgroup needs are fairly typical. You'll need file service, printer service, electronic mail service, and screen-sharing capabilities. It would be nice to provide room for ancillary networking products, such as meeting scheduling, when they become available. Each of the five members within your workgroup have been using a Macintosh for at least three months. Your workgroup's computing resources consist of the following:

- 1 Macintosh SE/30
- 2 Macintosh IICis
- 1 Macintosh IIfx
- 1 IBM PS/2
- 1 LaserWriter IINT

Additionally, your budget includes resources for adding two more workgroup members within the next 12 months, one of whom will be using a Macintosh IICi; the other a NeXT-cube.

Assessing Network Topologies

Expanding on the information presented earlier in this chapter, it's important to properly assess your present and future networking needs. Begin by evaluating the various network topologies in relation to your office's physical layout and your workgroup's networking needs.

- Bus topologies, such as the Macintosh's built-in AppleTalk capabilities, are inexpensive to implement and convenient for small networks. The major drawbacks to the bus topology is its limited size and total distance.

- Star topologies are convenient for use with in-wall telephone cabling, but require a star hub for every 10-30 users.
- Ring topologies are useful for mainframe connectivity and are quite common in PC networks. They require extra wiring for most in-wall cabling schemes, however.
- Internet backbones are used to connect several inter-networks and are useful for isolating network traffic.

The Macintosh's built-in AppleTalk bus topology is adequate for your needs. It's built-in, inexpensive, and easily implemented. Moreover, it's appropriate for your small network.

The next step is to assess the various types of cabling and cabling hardware that are available for your application.

Assessing Cabling Media and Cabling Hardware

The most basic part of the network is the *transceiver*. The transceiver is the hardware that sends and receives signals on the network. There are three main choices available to you: LocalTalk, Ethernet, and token ring.

- LocalTalk's single biggest advantage is that it's free. Its biggest drawback is that it's slower (230.4 Kbits per second) than the other options and inappropriate for installations that require heavy data traffic.
- Ethernet is the fastest of the three options available, with speeds of 10 Mbits per second. It also provides connectivity to both the VAX and the NeXTcube you'll be adding in the future. Its disadvantage is that it's significantly more expensive than LocalTalk.
- The token ring option offers the most consistent speed of the three options available to you and also offers IBM mainframe connectivity. Its drawbacks are that it is the

most expensive of the three options and that it's slower (4 Mbits per second) than Ethernet.

LocalTalk is the transceiver that is most appropriate for your short-term needs. In the future, especially when the NeXTcube is added to your workgroup, your needs will change. Your immediate needs are well met by LocalTalk because of its price (it's free on all of your Macintoshes), convenience, and availability for all devices.

Having decided to go with LocalTalk transceivers for your network, your next task is to evaluate the various media available for your application. Because LocalTalk transceivers are the most appropriate for your application, your media decisions are quite straightforward. At this time it is important to finalize the physical layout of your network, taking into consideration the location of various office layouts, existing wiring, and the like.

For the sake of brevity this discussion will assume that your office layout is quite typical, and is comprised of individual offices or individual office partitions for each workgroup member. This scheme tends to be the most common and the points discussed here are applicable to other office layouts as well. Figure 4.2 illustrates a typical office layout for a five-member workgroup with space already allocated for the two additional members that will be added within the next year.

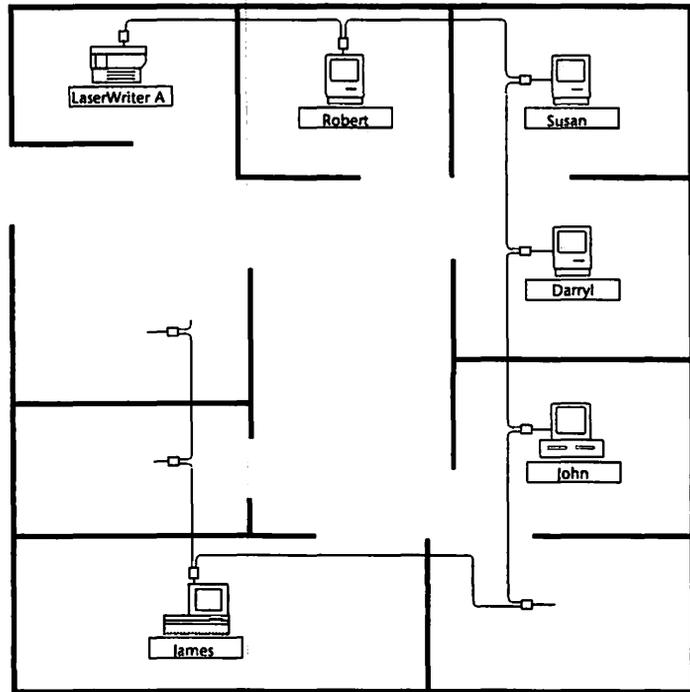


Figure 4.2 *Typical office layout for a small workgroup.*

Three basic media options are available to you: twisted-pair, coaxial cable, and optical fiber.

- Twisted-pair wiring is popular because of its low-cost and ease of installation. It is the same type of wiring used for standard telephone installations and its characteristics make it ideal for low speed networks of moderate length.
- Coaxial cable is ideal for those installations that need a higher level of performance than can be accomplished with twisted-pair wiring. It is widely used in networks like Ethernet because of its high bandwidth and excel-

lent shielding and is capable of speeds up to 16 million bits per second. Coaxial cable is the same wiring used for cable television.

- Optical fiber uses light instead of electricity to transmit the network signal. A laser generates up to a billion pulses of light per second which is transmitted through very fine glass fibers. This medium is very resistant to interference, is capable of extremely fast speeds, and is best used for installations that require long-distance, high-throughput connections.

Because you've identified your network services needs as file service, printer service, electronic mail service, and screen-sharing capabilities, standard twisted-pair LocalTalk wiring is the most appropriate media for your network. It offers low cost, moderate bandwidth, adequate length, simple installation, and high reliability. As your network grows, you can evaluate other options as it becomes necessary to do so.

You've carefully considered your group's potential needs in the future, however, and have decided to take a somewhat cautious route, leaving room for a possible Ethernet implementation in the future, especially considering the possible addition of a NeXTcube to your workgroup's computing resources.

Standard twisted-pair LocalTalk wiring will preclude attaching the NeXTcube to your network without adding either a bridge or Ethernet capabilities to all of the existing nodes. This is a fair trade-off for now, however. Standard twisted-pair LocalTalk is adequate for most small workgroups, and you don't really need Ethernet support for your existing needs. If it turns out that your network load is significantly greater than you had anticipated, you can always add Ethernet capabilities in the future.

The next step is to assess the network media available to you within the parameters you've defined for your network: LocalTalk twisted-pair wiring.

Assessing Twisted-Pair LocalTalk Cabling Systems

Since you've decided to use the Macintosh's built-in networking capability, your first task is to assess the various twisted-pair LocalTalk cabling systems available. You have two basic options at this point: Apple's standard LocalTalk cabling scheme or one of the third-party options that use standard telephone cable.

Apple's LocalTalk Cable System

Apple's standard LocalTalk cable will most likely be your initial choice, simply because you recognize that the LocalTalk network connection is built into each of your Macintoshes and you assume it's the best option available. It may not be.

Apple's LocalTalk cable system is designed as a standard workgroup cabling system that uses shielded twisted-pair wire. It's flexible and has set the standard for other LocalTalk cabling schemes. The LocalTalk cable system from Apple allows up to 32 devices to be connected per network segment and offers a bandwidth of 230.4 Kbps. A network that uses Apple's LocalTalk cable system can be up to 1,000 feet in length and segments of the network can be interconnected with bridges. The LocalTalk cable system from Apple comes in either 10- or 25-meter lengths and is available with either DIN-8 or DB-9 connectors. Prices begin at about \$75 per node.

Farallon Computing's PhoneNET System

Farallon's PhoneNET has been a viable alternative for LocalTalk cabling since shortly after Apple began shipping its LocalTalk cable system. What differentiates PhoneNET from Apple's LocalTalk cabling system is that PhoneNET employs ordinary telephone wire to implement cabling for LocalTalk.

The use of PhoneNET's standard telephone wire for LocalTalk offers three distinct advantages: lower cost, wider availability, and greater network distances. The Farallon PhoneNET system is capable of supporting network distances of up to 3,000 feet with ordinary telephone cable. In addition, the PhoneNET Repeater is available that doubles the effective network length of PhoneNET cabling by regenerating and reclocking LocalTalk signals.

Farallon's PhoneNET connectors are available in three models: DIN-8, DB-9, and DB-25. Prices begin at about \$60 per node with 10-packs available for about \$40 per node.

Farallon's PhoneNET system will also help you deal with the single existing IBM PC that exists within your workgroup's collective computing resources. In late 1990, Apple Computer turned to Farallon to develop and market AppleTalk products for IBM PCs and compatibles.

One of the first products to come out of the Apple-Farallon agreement was PhoneNET Talk, a \$195 software product that allows a standard IBM PC to access an AppleTalk network as a client. This software was originally marketed by Apple as AppleShare PC Client. The software will work with any Ethernet, token-ring, or LocalTalk card that supports the Open Data-Link Interface (ODLI) specification, and Farallon is also marketing a LocalTalk peripheral card for IBM PCs as a companion product. Named the PhoneNET Card PC•LocalTalk, the \$495 peripheral card enables a DOS or Windows PC to act as a client on the LocalTalk network.

Establishing Workgroup File Service

After assessing your workgroup's physical network needs, deciding on the physical layout of the network, and physically laying out and connecting the network, your first implementation task will likely be to establish file service for the workgroup.

File service consists of a file server (or servers) that make a volume accessible to multiple users on the network. It can be used as a central repository for the workgroup's data.

File servers can be used to make files and applications from your own computer available on the file server so that other workgroup members can have access to them. The server can also be used to access and use applications and files that are stored on the server just as if they were located on your own computer.

File servers are one of two distinct types: distributed or dedicated, although a network can employ both dedicated and distributed file servers.

Dedicated File Servers

Dedicated file servers require that a single workstation be dedicated to providing file service for the members of the workgroup, rendering the workstation used as the file server useless for other activities. A dedicated file server consists of the computer itself, specialized file server software (such as AppleShare, discussed later) and at least one high-capacity hard disk drive. A dedicated file server offers centralized administration, enhanced security, and the ability to assign and limit access privileges to specified users.

AppleShare

Apple Computer's AppleShare is the most popular dedicated file server for Macintosh networks. An AppleShare file server consists of a dedicated Macintosh running the AppleShare software. Any Macintosh from the Mac Plus on can be used as an AppleShare server. Likewise, any Macintosh from the Plus on can be used to access the AppleShare server.

The single biggest benefit to AppleShare is that it provides the best possible network performance. Because the AppleShare software runs on a dedicated server machine, it is less likely to crash because it is not susceptible to software anomalies inherent in application software or mistakes made by users.

Workgroup members that access the AppleShare server are referred to as clients. Access to the AppleShare server is gained through the familiar Finder interface and a Chooser device. Figure 4.3 shows the Chooser device used to log into the AppleShare server named "Savannah."

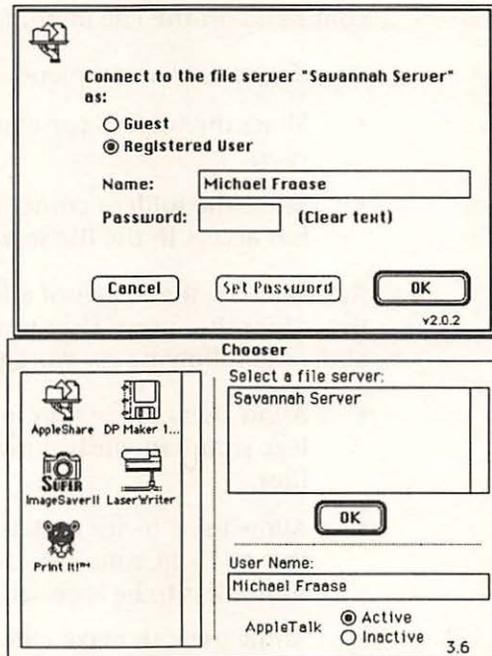


Figure 4.3 Logging into the AppleShare server named "Savannah."

After the log in sequence is successful, the AppleShare server volume appears on the desktop as any other volume except for a specialized icon indicating that it's a server volume. Once the server volume is mounted on the desktop, it can be used in the same manner as any volume that is attached to your local Macintosh.

The only differences between a mounted AppleShare volume and a local volume is that file operations are slower (about the same speed as a floppy disk) and each file and folder on the AppleShare volume is assigned a set of privileges. AppleShare privileges enable a folder's owner to specify certain attributes for the folder as well as the files contained within the folder. A folder's owner can use the Get Privileges command on the File menu to perform any of these tasks:

- Keep a folder completely private
- Share the folder's contents with a pre-defined group of users
- Make the folder's contents available to any user who has access to the file server.

Additionally, the owner of a folder can specify a set of privileges for other users. User privileges can be defined so as to allow the following for workgroup members:

- Allow users to see files in the folder. Setting this privilege simultaneously allows them to read and copy the files.
- Allow users to see folders inside of the folder. Setting this privilege simultaneously allows the contents of the subfolders to be accessed.
- Allow users to make changes inside the folder. Setting this privilege allows the users to create and modify files and subfolders within the specified folder.

The AppleShare file server software retails for about \$800.

Nondedicated File Servers

Nondedicated file servers are designed to run in the background on all of the computers connected to the network. The file service tasks are shared throughout the computing resources on the network and a dedicated server is not required. A nondedicated file server allows the workgroup members to selectively designate individual files and folders on their local hard disks for access by other members of the workgroup. The single biggest advantage of a nondedicated file server is the economy of eliminating the computer dedicated as the file server, allowing it to be used for other tasks.

A nondedicated file server is an appropriate choice for small workgroups that don't move a lot of information between workstations. A hidden advantage of a nondedicated file server is that you can always upgrade to a distributed or dedicated file server when the networking demands of your workgroup warrant.

Personal Server Network

Personal Server Network (PSN), from Calabasas, California-based Information Presentation Technologies Inc., is the only nondedicated file server for the Macintosh that provides a full set of AppleShare-compatible access privilege controls. Privileges are set using the Get Privileges command in the File menu from within the Finder or from the Access Privileges desk accessory within other applications.

Users access a PSN server through the standard AppleShare client software distributed as Macintosh system software. A Personal Server Network is the only nondedicated file server available that allows two servers, both using the AppleShare client software, to mount each other's disks.

The Personal Server Network software also allows each user to specify which of his or her hard disk drives to make avail-

able for file service use by other workgroup members. This is done by simply launching the PSN application, selecting the volumes to be published, and clicking the Publish button. Figure 4.4 shows two published volumes, named "Further" and "Intrepid."

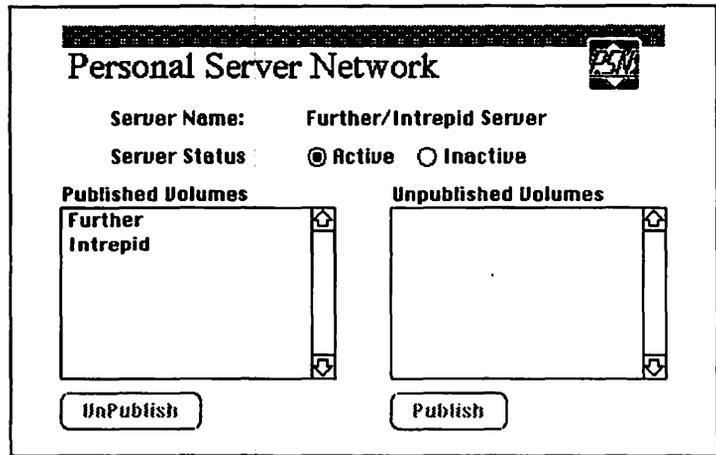


Figure 4.4 *The Personal Server Network published volumes dialog box.*

The procedure for unpublishing any volume is to select the desired volumes and clicking the Unpublish button. The published volumes dialog box is also used to quickly identify the server name and toggle the server status from active to inactive.

Each workgroup member that uses PSN to publish volumes for file sharing by other workgroup members configures his or her computer for use as a PSN server. Selecting the Server... command from the Configure menu and providing the administrator's name and password brings up the server configuration dialog box shown in Figure 4.5.

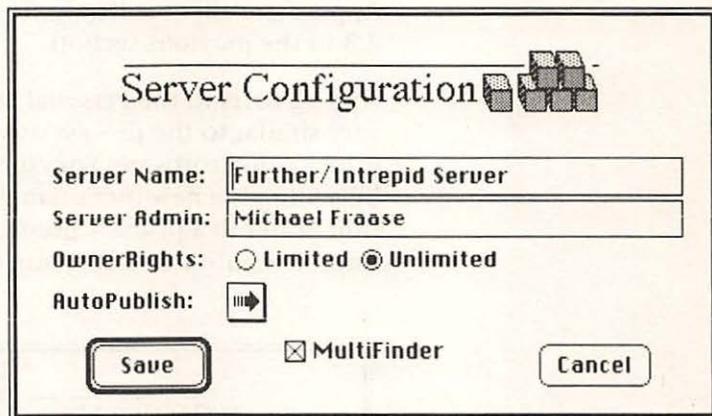


Figure 4.5 The Personal Server Network server configuration dialog box.

The server configuration dialog allows you to name the server, set owner rights, and specify the volumes to be autopublished when the application is launched. Any set of volumes can be autopublished each time the Macintosh is booted by setting the PSN application as the startup application.

PSN provides complete AppleTalk File Protocol (AFP) services for multi-user databases and also supports the Shared Environment feature. Using the Shared Environment, however, in limited mode slows performance for the local user by a factor of three or more.

While AppleShare's client software uses only about 25 KBytes of RAM, the amount of memory required by the PSN server software varies depending on the number of files and folders contained on the disk. Typical configurations require about 100 KBytes of RAM for each 80 MBytes of hard disk space. This sounds like a lot, but bear in mind that Macintoshes on the network that are used only as clients do not need to run the server part of the PSN package.

Users log into the PSN server volumes by using the standard AppleShare client software. The process is shown in Figure 4.3 in the previous section.

Adding users to the Personal Server Network file server is very similar to the process used under the AppleShare administration software. The administrator adds user accounts by adding the new user's name and password and assigning him or her to a primary group. All of these tasks are performed within the dialog box shown in Figure 4.6.

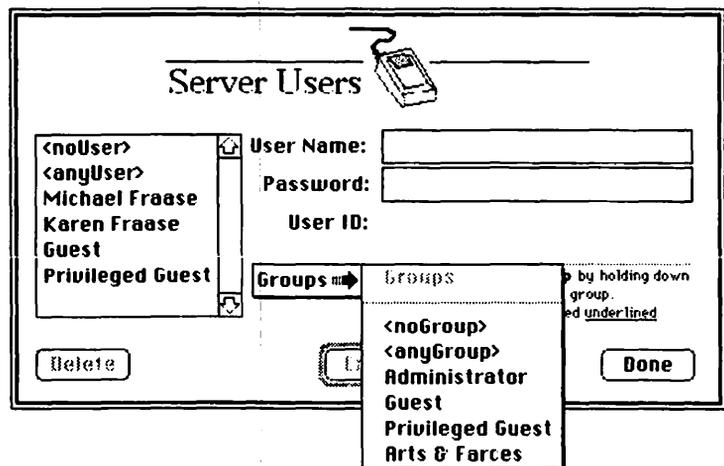


Figure 4.6 *The Personal Server Network user configuration dialog box.*

Personal Server Network, because it provides nondedicated file service, is about 30 percent slower than AppleShare running on a dedicated server Macintosh. It's a good choice for small workgroups that don't require a dedicated file server. Since it uses the standard AppleShare client software for client services, Personal Server Network file servers can also be accessed by IBM PCs or compatibles that are running AppleShare PC software.

The Personal Server Network software retails for about \$150 per client or server.

System 7.0 File Sharing

Built-in Macintosh file sharing is one of the most impressive features of System 7.0. The new file sharing feature, a superset of AppleShare, is built into the Finder. It supports all of AppleShare's features and adds the following:

- The user can specify selected folders or entire volumes to be published for file sharing use.
- A setting is available that allows the user to specify how much CPU processing time is available for file sharing tasks.
- Individual folders can be tagged to prevent their being moved, renamed, or deleted.
- Folders can be owned by groups as well as individual users. A folder's contents can be made available only to certain groups or individual users.
- Groups can contain subgroups.

The System 7.0 Finder has been updated to visually reflect the new file sharing services. A published folder appears on the desktop as a folder with attached network cables, and folders being accessed by network users are displayed with a user icon overlaid upon the folder's icon.

Opening the Users & Groups Control Panel displays a window with icons representing users and groups, as shown in Figure 4.7.

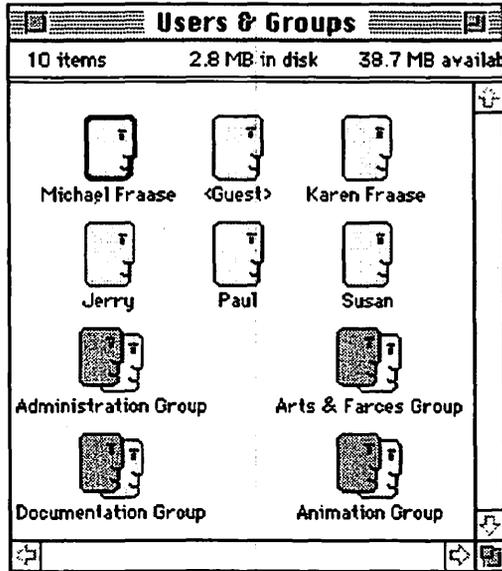


Figure 4.7 System 7.0 file sharing Users & Groups window.

The Users & Groups window displays all currently defined users and groups that have access to your Macintosh.

Double-clicking on a user's icon inside the Users & Groups window allows you to define a password and specify access privileges for that user's access to your Macintosh. Figure 4.8 illustrates a sample user's access privileges. In this case, the user selected is Karen Fraase, represented by the icon in the upper right portion of the Users & Groups window.

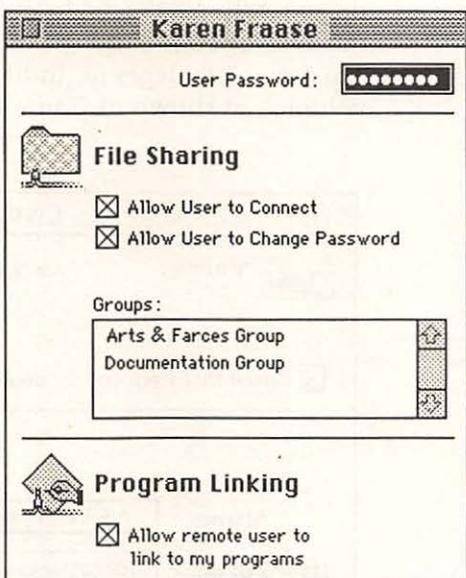


Figure 4.8 System 7.0 file sharing user's access privileges.

The user's access privilege window lets you set the file sharing settings for each user. Checkboxes are provided that allow the user to connect to your Macintosh and change his or her password. Unchecking the Allow User to Change Password option is useful when you define guest accounts, for example. Unchecking the option would prevent the user from changing the password associated with that user account.

Any groups that the user belongs to are displayed in the scrolling list in the center of the user access privilege window.

At the bottom of the window, the Program Linking option allows remote users to forge links to applications stored on your hard disk drive. This feature is discussed later in this section.

Groups can be defined and users can be added to groups by dragging the user's icon to the group icon. In addition, using the Sharing command within the Finder lets you set up special access privileges for individual documents and entire volumes, as shown in Figure 4.9.

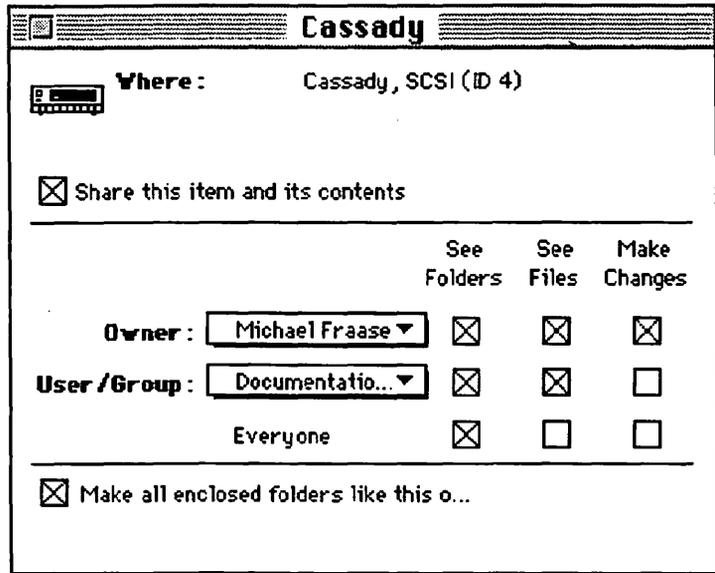


Figure 4.9 System 7.0 file sharing volume access privileges.

Apple suggests that a maximum of 50 users and groups be defined for each Macintosh and only 10 users may be simultaneously logged into your Macintosh. A slider control is available in the Sharing Setup Control Panel that lets you define the priority level for file sharing on your Macintosh.

Apple's System 7.0 also includes a Publish and Subscribe feature, also referred to as live copy and paste. This is the feature that was referenced earlier in the user access privilege window and the Program Linking setting is shown in Figure 4.8. Publish and Subscribe allows information contained in

one document to be copied into another document along with a live link to the originating document. When the original document is changed, the linked document is updated automatically. If the linked document is open at the time the original is changed, the update takes place immediately. If the subscribing document is not open, it is automatically updated the next time it is opened.

The Publish and Subscribe process is controlled by the Edition Manager. New commands on the Edit menu of applications that support this feature of the System 7.0 operating system will let users publish information, making it available to other documents and users. A subscribe command will allow information to be imported that remains linked to the originating document. More than a single document can subscribe to a single originating document and subscriptions work across a network. Commands will also be available that will allow you to freeze all links, stopping subsequent updating.

The Publish and Subscribe feature of System 7.0 is best thought of as a one-way communication channel. Documents can subscribe to an originating document that has been published, but there is no communication from the subscriber to the publisher.

File sharing service under System 7.0 also supports Interapplication Communications over the network. This allows you to configure one Macintosh to accept instructions over the network from other workgroup members. This can be used, for instance, to run applications on a remote machine and deliver the results back across the network.

Interapplication Communication features of System 7.0 are controlled by AppleEvents. AppleEvents, in contrast to the Edition Manager's Publish and Subscribe functionality, is best thought of as a two-way communication channel. This two-way communication channel allows documents to send commands, instructions, results, and answers between themselves.

Distributed File Servers

DataClub

International Business Software Inc., of Sunnyvale California markets DataClub, the only distributed file server for the Macintosh. As the first truly distributed file server, DataClub shares the task of file service across multiple computers on the network. Users, however, see a single icon on the desktop that represents the distributed server.

DataClub is AppleTalk File Protocol (AFP)-compliant, so any software that works with AppleShare will work with DataClub. This includes the standard AppleShare client software as well as AppleShare PC for IBM PCs and compatibles.

The DataClub server software also emulates all the security features of AppleShare, including the byte-range locking feature of AFP's Shared Environment.

The software runs on any Macintosh from the Plus or greater, and requires about 450 KBytes of RAM for server activities. While this appears to be a disproportionate amount of memory, remember that workgroup members that will not be providing distributed file service to other members of the workgroup need only run the AppleShare client software, requiring only about 25 KBytes of RAM.

To add a DataClub server to the network, you simply hook the Macintosh into the network and install the DataClub software. The new server automatically exchanges information with all the existing servers and updates all user and group information without user intervention.

Installing the DataClub software is accomplished by dragging the DataClub startup document to the System Folder of each Macintosh on the network that will be used as a file server and restarting the computer.

When the Macintosh is restarted, the DataClub startup document is accessible via a Control Panel device, where the basic server functions are defined. Figure 4.10 shows the configuration options available within the DataClub Control Panel device.

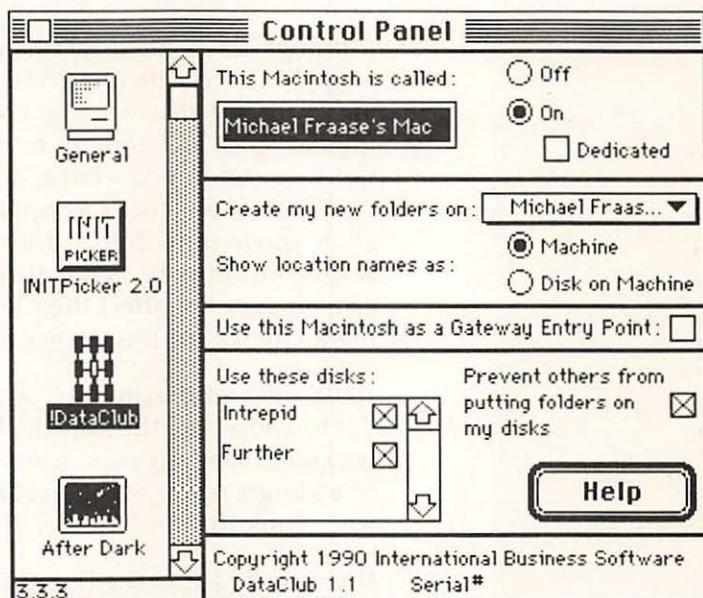


Figure 4.10 The DataClub Control Panel configuration options.

Controls are available to let you turn file service on and off at each local Macintosh server, as well as a dedicated setting that, when selected, allows the Macintosh to be used as a dedicated DataClub server.

A check box is also provided for preventing other users from putting their folders on your hard disk drives used by the DataClub virtual server. This is one of the best features of DataClub because it allows you to add disks to the total

amount of server space while at the same time restricting other users from putting their data on your physical disk. This arrangement lets you share files that reside on your physical disk with workgroup members without giving up any of your physical storage space. Additionally, the DataClub software allows you to define how much CPU time is dedicated to file service activities.

Although DataClub displays only a single file server icon to logged-in users, the physical files are actually stored on various hard disk drives anywhere on the network. A virtual file's physical location can be obtained by using the Finder's Get Info window. On DataClub servers, the location information field is actually a pop-up menu with commands that allow you to move folders from one physical disk to another across the network. While this action moves the files physically, it does not affect their logical placement in the folder hierarchy within the server's file structure.

DataClub's administration program, Admin, is visually similar to and works the same as AppleShare's administration program. The only significant difference is that the process of adding a server with DataClub is much easier than it is with AppleShare.

Users are added to the DataClub server by selecting the Create User command from the Users menu. This action brings up DataClub's Define User dialog box, shown in Figure 4.11.

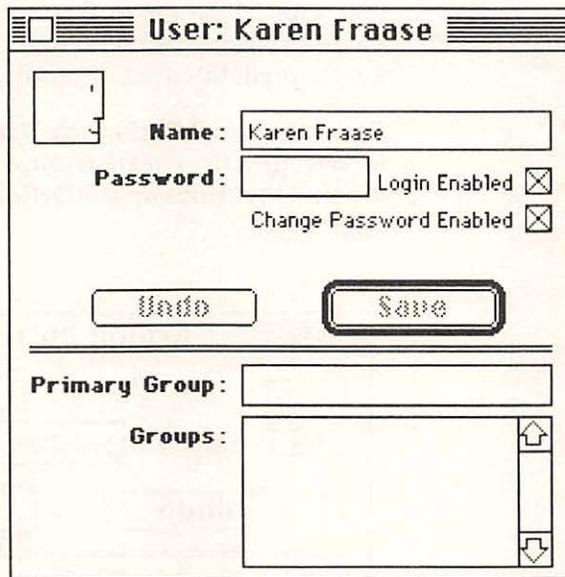


Figure 4.11 The DataClub Define User dialog box.

The Define User dialog box allows you to define the name and password for each user. Checkboxes are also provided that let you enable login privileges and specify whether or not the user can change the password assigned to the account. This latter feature is useful for defining guest accounts or other login accounts that may be shared by several users.

Provisions are also provided in the Define User dialog box to assign the user to a primary group.

Groups are used to predefine categories of users for access privileges to a common group of folders and files on the virtual server. A group is usually comprised of a subgroup of the general workgroup. A product development workgroup, for example, may be comprised of development and marketing subgroups.

Defining a group within the file server environment allows you to provide new workgroup members with a basic set of access privileges by including them in one of the groups that have a predefined set of privileges.

Defining groups within the DataClub Admin utility is done by selecting the Create Group command from the Groups menu. This brings up the Define Group dialog box shown in Figure 4.12.

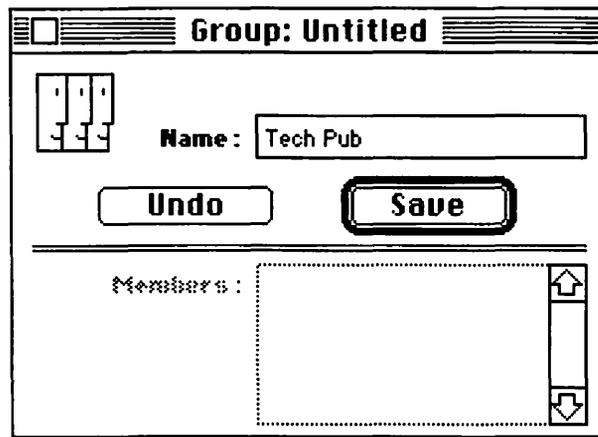


Figure 4.12 The DataClub Define Group dialog box.

One of the most significant administration advantages DataClub enjoys over AppleShare or Personal Server Network is that it's virtually impossible to lose the user, group, and privilege information for the workgroup. This information is automatically retrievable in the event of a crash or a problem unless every DataClub server on the network is destroyed at the same time. As long as a single server Macintosh survives any network-wide problem, other Macintoshes on the network will automatically be updated when they are brought up and the DataClub software is restarted.

Error and crash handling within the DataClub server environment is also more seamless and fault-tolerant than AppleShare. If a single DataClub server node crashes, the only nodes affected are those that are using files that are physically stored on the crashed server's disk. Other nodes wouldn't even notice that one of the virtual server computers crashed until they tried to access a file that was physically stored on the crashed Macintosh.

DataClub can also be configured in a manner that would allow it to perform at a faster speed than AppleShare over identical network media. It's possible to set up the DataClub virtual server on a series of dedicated Macintoshes. Because the file service tasks are distributed across the dedicated Macintoshes, the workload can be balanced between the computer power offered by the bank of server computers.

The DataClub server software retails for about \$300 for three server nodes and about \$800 for ten server nodes.

Assessing File Service Options

The perfect file server for all workgroups simply doesn't exist. There is not a single software package that can address all the needs of every workgroup configuration. The best place to start in analyzing your choices for file service, then, is to assess your workgroup based on the following three criteria:

- physical size of the network in number of nodes
- size of the files to be transferred across the network
- future growth of the number of nodes on the network and the size of files likely to be shared.

Obviously, the number of members of your workgroup will tend to limit the options available to you. If your workgroup is larger than five members or so, or moves large amounts of

data across the network, you can safely rule out any nondedicated file server option.

For small workgroups, Personal Server Network is a viable option. If you anticipate that your workgroup will expand in the near-term, however, it may be an inappropriate choice since you will need to move to a dedicated or distributed server and your investment will be lost. In that case you'd be better served by the DataClub distributed server.

DataClub is a good option for small workgroups that move a lot of information back and forth across the network as well as medium-sized workgroups of up to 10 or 12 people. In the event that your workgroup—or the size of the information it moves across the network—grows, you can always use the DataClub software to provide distributed file service that offers a higher performance level than even a dedicated file server.

AppleShare is most appropriate for large workgroups of more than 15 or so members. It's also the best choice for enterprise-wide file service between workgroups within a large organization. It's the standard and may already be in use within your organization. A dedicated AppleShare server is also the best option if your workgroup is fairly sizeable and plans to implement an electronic mail system. Both of the leading electronic mail servers for the Macintosh can run as a companion task on the same Macintosh used for the AppleShare file server.

For those workgroups already running System 7.0, the built-in file sharing features of the new operating system may prove to be perfectly adequate—especially for smaller workgroups. It certainly won't cost you additional money to find out for yourself.

Additionally, System 7.0's built-in file sharing capabilities can be used on networks with Macintoshes running version 6.0x system software. Workgroup members using System 6.0x can access the file sharing features of other workgroup

members through the use of the standard AppleShare client software.

For workgroups of small to average size—no more than five or six members—sharing information of average size, the built-in file sharing inherent in System 7.0 is the best choice.

Establishing Workgroup Printer Service

As discussed earlier, a print server is a hardware- or software-based network service that lets multiple users on the network send files to the same printer at the same time. This allows the workgroup members to share expensive printers.

The printer stores the files in a queue and processes them in the order in which they were received. Some print servers allow an administrator to override the order in which the files are processed.

The main advantage of a print server is that it allows a user to return to productive work after sending the print job to the printer rather than waiting for the printer to finish its job.

Two common alternatives for workgroups are to use a print spooler or a print server. An example of each alternative is discussed below.

MultiFinder and System 7.0 Print Spooling

Automatic print spooling is built into the Macintosh operating system when used under MultiFinder or System 7.0. The Macintosh operating system's print spooling features are best described by breaking the actual process down into its components.

- *Print spooling* is the process of writing a representation of a document's printed image to disk and feeding the disk file to the printer as it can take it. The advantage to print spooling is that it frees up the computer for other tasks, allowing you to get back to work quicker.
- A *print spooler*, such as the one built into MultiFinder and System 7.0, is a utility that writes the representation of a document's printed image to disk, schedules it to print in a queue of other jobs, and then prints it.
- The *print queue* is a collection of spooled documents, awaiting printing, that are stored on the print server disk and printed in order.
- A *print server* is a combination of software and hardware that stores documents sent to it over the AppleTalk network and manages the printing of those documents on a LaserWriter or other network printer.
- The *Print Monitor* is an application that monitors background printing and provides options intended to give you additional control over what happens to documents you are printing.

Under MultiFinder and System 7.0, you can print to a LaserWriter while continuing other work by using the operating system's built-in background printing feature. Turning on the background printing feature is a simple matter of selecting the Background Printing radio button in the Chooser as shown in Figure 4.13. Under System Software versions prior to 7.0, MultiFinder must be active to turn background printing on.

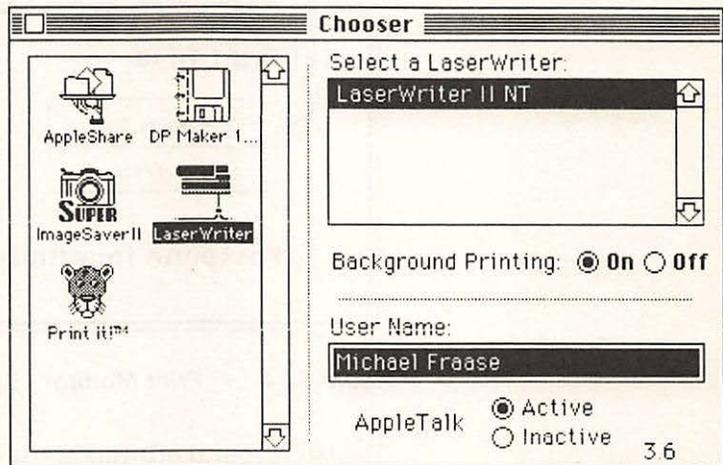


Figure 4.13 Turning on background printing in the Chooser.

Once the background printing selection has been turned on, print your documents as you ordinarily would to any device on the network. The documents will automatically be spooled to the Spool Folder inside your System Folder.

A companion application, Print Monitor, is launched automatically when the spooling process begins. Print Monitor resides in your System Folder and can be used to manipulate the order in which documents are printed. This reordering is limited only to your local documents and has no effect on documents printed by other workgroup members. Figure 4.14 shows the options available to you for manipulating the order in which your documents are printed.

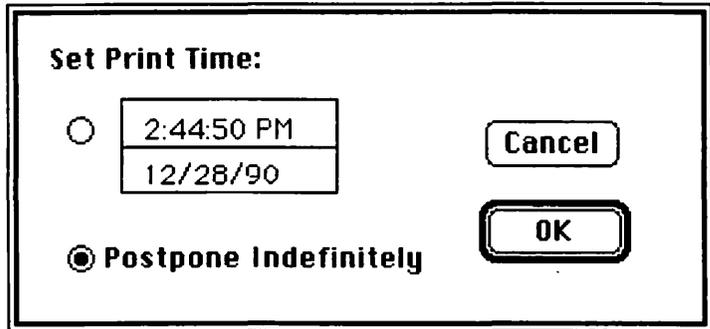


Figure 4.14 Print Monitor's printing schedule options.

The background printing feature of the Macintosh operating system provides the added benefit of built-in error recovery. If, for example, you have a system error after your document has been spooled, simply restart your Macintosh. The system will ask you if you want to continue to print the documents that have been spooled to disk but not printed.

AppleShare Print Server

Apple Computer's AppleShare Print Server allows all the members of a workgroup that are attached to the network share a spooling facility for LaserWriter printers. The AppleShare Print Server is designed to run as a companion task on a Macintosh used as a dedicated AppleShare file server. Alternatively, if file service is not required, the AppleShare Print Server software can be installed on a dedicated Macintosh to provide print service to the workgroup.

The AppleShare Print Server enables workgroup members to regain immediate use of their computers by transmitting the document for printing to a centralized server. In addition, the order of printing for queued documents can be changed,

allowing high priority documents to be printed before other documents.

Because the AppleShare Print Server is a server, any number of users can simultaneously send documents to be printed. The server performs administration duties and manages the print queue automatically. Also, since the AppleShare Print Server appears on the network as a LaserWriter, the print server can be accessed by any program or device on the network, including LocalTalk-equipped IBM PCs or compatibles. Any LaserWriter-compatible printer will work with the AppleShare Print Server, including the Varityper and Linotronic families.

According to Apple, users can save more than 70 percent of normal printing time by using the AppleShare Print Server because the document printing processing time and the queue time are both passed through to the server.

Establishing Electronic Mail Services

Electronic mail is a software-based network service that allows workgroup members to communicate with each other by exchanging messages that may also include attached files. Email is the groupware equivalent of the interoffice memorandum, although it offers significant advantages over the traditional memo. Electronic mail is delivered instantaneously, for example, and does not require the use of paper. Additionally, some Email systems allow workgroup members to communicate with people outside of the local workgroup, using telecommunications services like AppleLink and CompuServe.

When a workgroup member sends an Email message to another user, the message (and any attached documents) is stored on a central mail server. The recipient is notified that he or she has mail waiting and can retrieve it the next time they log into the mail server. Most mail servers do not re-

quire a dedicated Macintosh although some can be used in conjunction with dedicated file servers.

Electronic mail is very efficient because the recipient doesn't have to be available—their computer doesn't even have to be turned on—at the time the message is sent. Additionally, a single message can be sent to various recipients and need only be sent once.

There are two general strategies of electronic mail service:

- *Point-to-point* electronic mail allows one node on the network to send a file directly to another node.
- *Store-and-forward* electronic mail systems use a central mail server to store all messages and files. The files and messages are sent to the recipient when they are requested.

The following sections will focus on the two most popular store-and-forward electronic mail systems for the Macintosh. Point-to-point file transfer utilities are covered later in this chapter because they are appropriate for electronic mail service only in the smallest of workgroups and only on the smallest of networks.

Workgroup Electronic Mail Considerations

An electronic mail system can be judged adequate for workgroups providing it meets the following criteria:

- *Enterprise-wide connectivity.* Every member of the enterprise should be able to communicate with any other member, even across various hardware platforms.
- *Multiple message systems.* Dividing the enterprise's electronic mail needs into smaller components improves efficiency. This requires the support of multiple message bases.

- *Support for large numbers of users.* An appropriate electronic mail system includes room for future expansion.
- *Store-and-forward messaging.* Electronic mail messages (and any attached files) are sent to a mail server rather than an individual computer. This allows workgroup members to send email to their associates without regard to whether or not the recipient's computer is even turned on.
- *Automatic notification of delivery problems.* If an electronic mail message cannot be delivered within a predetermined time frame, it is crucial that the mail system notify the sender and return the undelivered messages.
- *Transparency.* A user should not be required to know anything more than the recipient's name in order to send electronic mail to anyone within the enterprise. The various mail servers contained within the electronic mail system should automatically configure themselves when users are added or deleted. This allows workgroup members to send messages among themselves without knowing the physical location of the recipient.
- *Integration with applications.* An appropriate electronic mail system must be integrated as seamlessly as possible into the computing environment. Ideally, workgroup members should be able to exchange information from within any application.
- *Gateways to other message systems.* Large organizations frequently use more than one electronic mail system. An appropriate mail system must be able to communicate with a wide variety of other messaging systems. Because workgroup members may be outside of the main organization—vendors and consultants, for example—it's important that the electronic mail system be capable of allowing communications via public and telecommunications systems such as CompuServe, AppleLink, and MCI Mail. These external gateways

must be as flexible as possible, allowing the network administrator to customize communication parameters, account information, and connect times.

This set of parameters is met by the two leading electronic mail systems for the Macintosh, CE Software's QuickMail and Microsoft Corp.'s Microsoft Mail. The following sections provide overviews of both products.

QuickMail v2.2.3

QuickMail v2.2.3, from West Des Moines, Iowa-based CE Software, is a complete desktop communications system that enables workgroups to complete three basic collaborative computing functions:

- Send and receive electronic mail
- Send and receive files
- Real-time conferencing

In small workgroups of less than five members, QuickMail may provide all the basic networking tools necessary for collaboration. Most workgroups, however, will want to supplement QuickMail with other collaborative computing tools.

QuickMail Bridges

One of QuickMail's strongest points is that it offers better bridging functions between disparate mail systems than any other product on the market. The level of transparency of this bridging functionality, from the workgroup member's perspective, is exceptionally high. With QuickMail, workgroup members never have to think about *where* the mail is going, only to *whom*.

Figure 4.15 shows a typical QuickMail address book that is comprised of addresses on multiple mail systems, specifically a local QuickMail server (referred to as a MailCenter in QuickMail terminology), a printer bridge, a CompuServe bridge, and an AppleLink bridge.

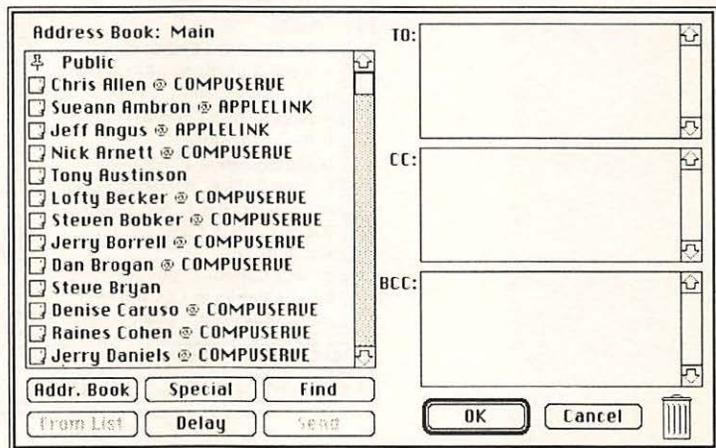


Figure 4.15 QuickMail address book for multiple mail systems.

A variety of bridges to public and private mail systems is provided with the QuickMail distribution package. Additionally, bridges to just about any other mail system are available from third-party developers. The underlying strength of the mail bridges available for QuickMail is that they do not require the workgroup member to know anything at all about the system being bridged to. All he or she has to do is select the address from the QuickMail address book.

Using the address book shown in Figure 4.15, for example, I could send the same message to Denise Caruso, Nick Arnett, Jeff Angus, and Steve Bryan even though the addresses are on three different mail services. The procedure is as simple as dragging the addresses in the left panel of the dialog box

to the appropriate panel in the right side of the dialog box. Sending a message to Jeff Angus with carbon copies to Denise Caruso and Nick Arnett and a blind carbon copy to Steve Bryan is shown in Figure 4.16.

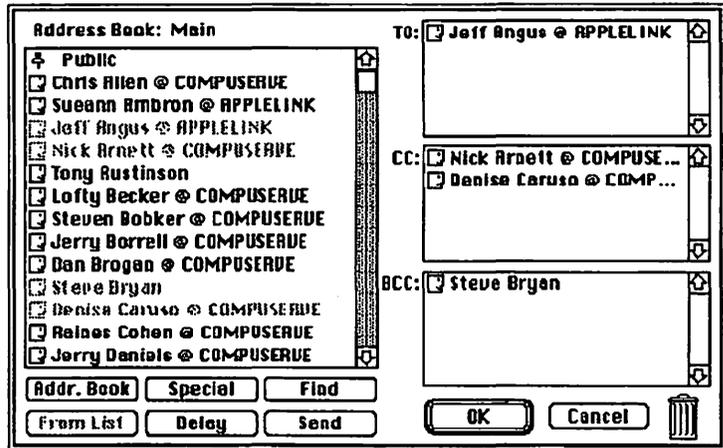


Figure 4.16 Addressing a mail message to four recipients on various mail systems.

Bridges are provided for most of the national telecommunications services. A QM-QM bridge is also provided, allowing two QuickMail servers to connect via modem and exchange mail. This is an important feature for workgroups with more than one location, as well as for intra-collaboration between various workgroups.

Adding bridges is done by dragging the appropriate file to the QA folder inside the QuickMail folder and configuring the new bridge from within the QuickMail Administrator application. Figure 4.17 shows several bridges within the QuickMail Administrator application.

QM Administrator					
Status idle				QM Server	
Administrator:		Time: 2:11 PM	Disk Space: 33553K	NameServer	
Modem: No modem		Memory: 465K	Dead Mail: 0	50 of 1000 names	
MailCenter	Type	Next Connection	Waiting	Urgent	Status
APPLELINK	QM-Link		0	0	
COMPUSERVE	Telecom		0	0	
FARCE_MAIL	Online				
GENIE	Telecom		0	0	
QMCONCIERGE	QMConcierge™	2/6/91 2:07 PM	0	0	
LUUCP_BRIDGE	UMCP\QM™	2/6/91 2:07 PM	0	0	

Figure 4.17 Bridges within the QuickMail Administrator application.

You configure any of the bridges by double-clicking on it and filling in a series of dialogs related to account identification, user definitions, connection schedules, connection script, and other information related to establishing the connection.

Double-clicking on the CompuServe Telecom bridge, for example, brings up the dialog box shown in Figure 4.18. This dialog is used to assign a password for access, a custodian to receive log files of bridge activity, and a specification for how often to send the log to the custodian.

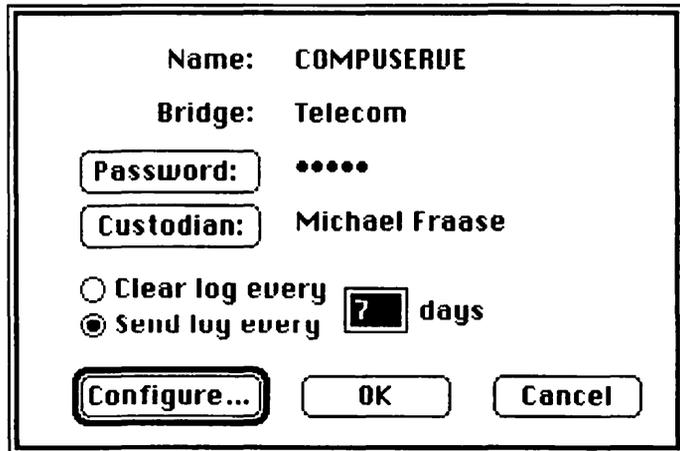


Figure 4.18 *QuickMail Administrator's initial configuration dialog box for the CompuServe telecom bridge.*

Clicking the **Configure...** button brings up a nested dialog box, shown in Figure 4.19, that allows you to specify communications parameters for the connection. This dialog box allows you to specify the phone number of the bridge, the baud rate for the connection, your account and password on the commercial service, and assign a script to perform when the connection is established.

The script is an editable text file that allows you to finely configure the actual communications link. The script for exchanging mail with CompuServe, for example is quite simple, and it is possible to create scripts for other services.

Telecommunications Bridge for QuickMail™ v2.2.3

Phone#:

Baud: 1200 2400 4800 9600 19200

QuickMail Network
MailCenter Name:

Commercial Service

Account #:

Password:

CompuServe.Script.v2.0

Display connection

Include headers

Figure 4.19 QuickMail Administrator's communications settings for the CompuServe bridge.

Clicking the Set Connect Times button brings up yet another dialog that lets you specify the times to attempt to establish a connection with the remote service. The connect times dialog is shown in Figure 4.20.

QuickMail provides extensive customizability for defining connect times. The settings shown in Figure 4.20 call for QuickMail to automatically dial CompuServe when 10 or more messages are waiting to be delivered or when 1 or more urgent messages are waiting. (QuickMail allows each workgroup member to assign a priority level to each message he or she sends.)

Additional controls are provided for establishing connections with the remote service on a recurring schedule, every 15 minutes during normal business hours, for example.

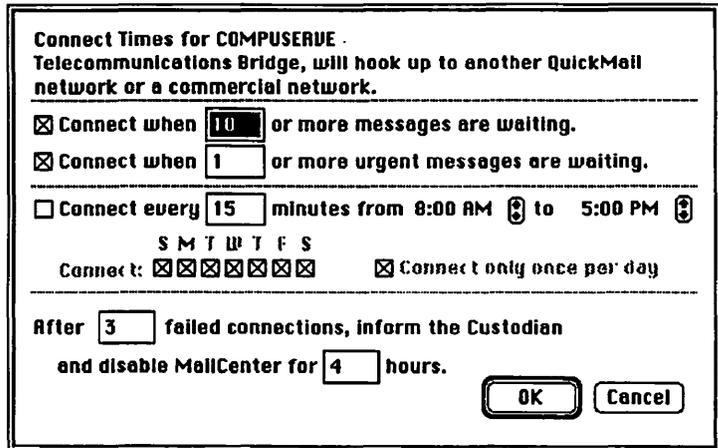


Figure 4.20 QuickMail Administrator's connect times dialog box.

Remote Access

Individual workgroup members that are not physically connected to the local area network at the organization's central facility can still exchange electronic mail with other workgroup members. A stand-alone application, QuickMail Remote, is included in the QuickMail package for geographically dispersed workgroup members. QuickMail Remote allows a workgroup member to log into the QuickMail server and exchange mail via modem from a remote location.

If the workgroup is comprised of several subgroups, each of which is geographically dispersed, it's a simple matter to set up a bridge that links the QuickMail servers at the various sites throughout the day.

For more than two remote workgroups, a star configuration can be setup with a single QuickMail server acting as a stor-

age terminal for all remote systems. This lets you design a system that requires each remote system to call only the central location than each remote site.

All that's needed to implement such a system would be to create a Telecom MailCenter for each remote site on the central QuickMail server. Each remote site would create a Telecom MailCenter for the central server. All sites would configure the MailCenters for dialup service at prearranged times.

Bulletin Board

A public bulletin board is also available for messages intended for wide distribution on the local area network. This feature is more useful than it would first appear because the addressing sequence is streamlined by only having to select the Public address. Also, only a single message is created on the server, reducing disk fragmentation and enhancing performance.

This feature allows small workgroups to enjoy a sort of pseudo-file-service on those networks without a file server. Software updates, for example can be sent as file attachments to a mail message posted to the public bulletin board. Each workgroup member can retrieve the new software at his or her convenience.

Groups

It would be ideal if QuickMail provided complete bulletin board functionality with multiple message sections. This would allow workgroups to maintain ongoing discussion groups within the mail system. The closest thing to this is the ability to assign a group or series of groups for use in sending mail messages and file attachments.

For example, you may define a QuickMail group that corresponds with the groups you have defined on your file server. A technical publications workgroup, for example, could create a QuickMail group that was comprised of all the members as well as a series of subgroups. The writers in the group could form a subgroup, the illustrators another subgroup, and so on.

This type of setup would allow an editor to send a message to all the writers in the workgroup concerning a grammatical style change, for example. Such a message would be composed and sent to the Writers group rather than each individual writer.

This series of groups and subgroups would streamline the information flow within the workgroup by allowing multiple recipients to be addressed as a group rather than individually. It would also serve to enhance the productivity of the workgroup in general since information of a grammatical style change would be of no importance to the illustrator members of the workgroup, and the illustrator group would be spared the flow of messages that were of no importance to them.

File attachments to mail messages are another strong point of the QuickMail system. Up to 16 files can be attached to each QuickMail message. The attached files can be virtually anything you can do on your Macintosh: Graphics, animation, sound, spreadsheet information, databases, etc.

Mail Notification

Each workgroup member connected to the local area network that contains a QuickMail server can select between three ways of being notified that mail addressed to them has been delivered to the server:

- *Visually.* A flashing QuickMail icon covers the Apple menu item in the menu bar when new mail is received.
- *Auditory.* The Macintosh plays a sound when new incoming mail is received.
- *Both.* The icon flashes and a sound is played when new mail is received.

Mail is received, created, and sent in the QuickMail system by selecting the QuickMail desk accessory. This brings up a window like the one shown in Figure 4.21 that displays all of your mail. Unread mail is displayed with an unshaded icon. Messages that have files attached are displayed with a miniature appendage.

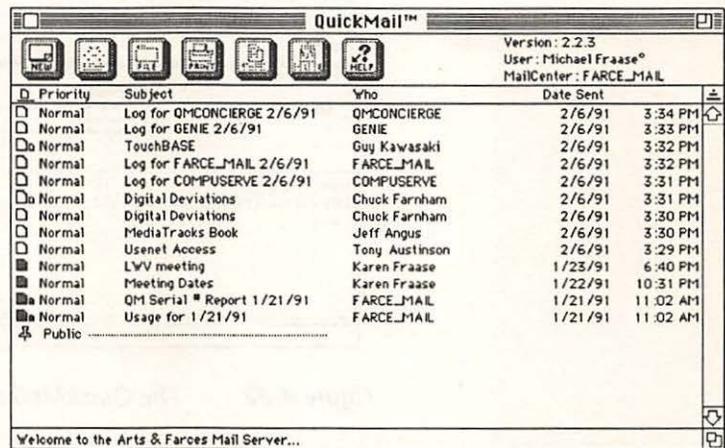


Figure 4.21 The QuickMail main window.

Mail is read by double-clicking on one of the message items or by clicking on the message in the list to select it and then clicking on the Read button. The message is then opened in another window, like the one shown in Figure 4.22.

A range of options, represented by the row of icons across the top of the window are provided for taking action on the message. Options are provided for filing, printing, saving as a text file, opening any enclosures, turning on the receipt function, forwarding the message to someone else, and replying to the message.

Additionally, the button labelled REC allows you to record a voice message that is attached to your text message if you have a newer model Macintosh or a sound digitizer on an older Macintosh.

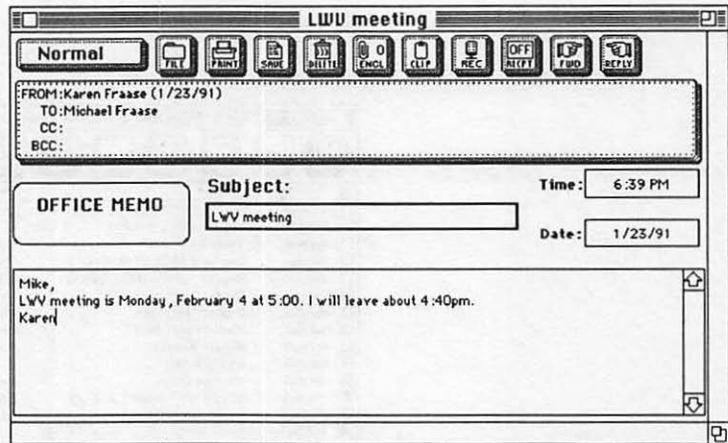


Figure 4.22 *The QuickMail message window.*

QuickMail messages that you wish to keep for future reference can be organized into folders that appear and operate like standard Macintosh file folders on the desktop.

Conferencing

A real-time conferencing feature is also built into the QuickMail system. The conferencing feature lets you type messag-

es over the network to any number of recipients at the same time. It's most useful for short, information-dense messaging that needs to be done in real time. Something that would ordinarily take a very short telephone call—scheduling a small meeting, for example—is a good candidate for QuickMail's conferencing feature.

The conferencing feature is unique among electronic mail products for the Macintosh and is much more useful than one would imagine. An example of QuickMail's conferencing feature is shown in Figure 4.23.

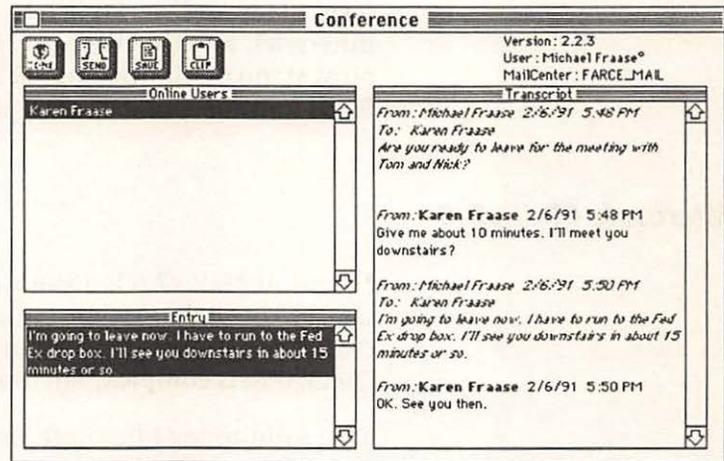


Figure 4.23 The QuickMail Conference window.

The complete transcript of the real-time conference can be saved to disk, copied to the Clipboard, or printed.

A conference is initiated by selecting the Conference command from the QuickMail menu. This opens the Conference window with a listing of everyone available on the network. To send a message, select the recipient or recipients, type your message in the Entry panel, and click the Send button.

The QuickMail Conference window will automatically be opened on the recipient's Mac displaying your message.

If your network is set up with different zones, you can use the Zone button to specify the zones to use for conferencing.

A privacy feature is built into QuickMail that prevents other workgroup members from disturbing you with conference requests for a period of time you specify.

CE Software also provides a specialized forms customization utility, QM Forms, that lets you design forms for special purposes. Forms can be designed for "while you were out" messages, daily and weekly schedules, routing requests, travel itineraries, and the like. QM Forms is designed to work like most standard drawing applications and has predefined fields for time and date stamping, subject categories, etc.

Microsoft Mail v2.0

Microsoft Mail v2.0 is a study of contrasts compared with QuickMail v2.2.3. Where QuickMail is rich and deep, Microsoft's product is spartan and relatively shallow. Where QuickMail is complex, Microsoft Mail is simple.

That's not to say Microsoft Mail is not a good product; it's just that it approaches electronic mail from a completely different perspective than QuickMail does. Microsoft Mail is designed to provide basic electronic mail service; nothing more.

Microsoft Mail is setup in a manner similar to QuickMail. A startup document (INIT) is placed in the System Folder of the Macintosh designated to be the mail server and an INIT and Desk Accessory are installed on all client Macintoshes.

Server Administration

A Microsoft Mail system is markedly simpler to administrate than a QuickMail system, mostly because there's less to administrate. Where QuickMail uses a separate application for defining users and configuration tasks, Microsoft Mail builds it into the mail system itself. To administrate the Microsoft Mail system, the network administrator simply logs into the server as the Network Manager.

User accounts are defined and modified by selecting the Users and Groups command from the Mail menu after logging into the server as the Network Manager. Figure 4.24 shows Microsoft Mail's Users and Groups window

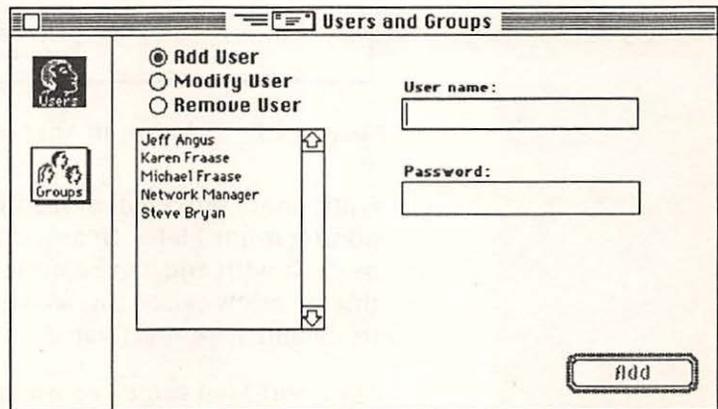


Figure 4.24 Microsoft Mail Users and Groups window.

Where QuickMail allows each user to define and configure his or her own groups, Microsoft Mail requires the network manager to define all groups. This is severely limiting in certain circumstances, and a fairly serious design flaw for autonomous workgroups. It works well, however, in large organizations that are highly structured in nature.

Mail Notification

Microsoft Mail users are notified of waiting mail by a dialog that pops up when new mail is received, as shown in Figure 4.25.



Figure 4.25 Microsoft Mail notifier window.

Options are provided for reading the new mail now, or postponing it until later. In any case, the new mail dialog must be dealt with and can be quite disruptive at times. This notifier window can be disabled in the Preferences settings, but its default state is activated.

Microsoft Mail can be configured to automatically log the user into a selected mail server. When Microsoft Mail is selected from the Apple menu, the user is prompted for his or her name and password. The server is selected via the Chooser.

When the workgroup member is properly logged into the mail server and Microsoft Mail is selected from the Apple menu, Microsoft Mail's main window is displayed listing all existing mail. A sample is shown in Figure 4.26.

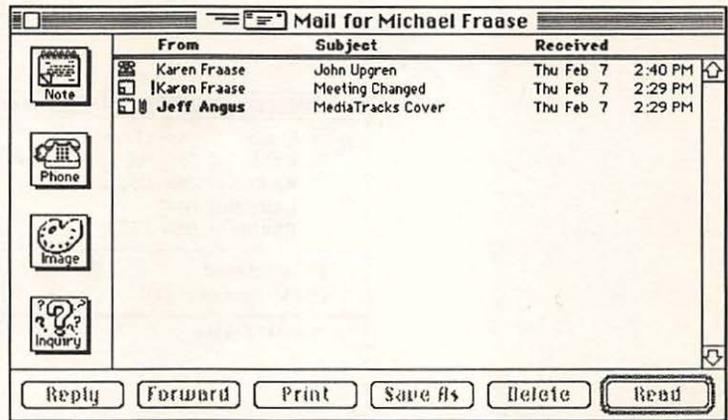


Figure 4.26 Microsoft Mail main window.

Clicking any of the four icons along the left side of the window opens a new mail message of the selected type: Note, Phone, Image, or Inquiry.

Buttons along the bottom of the window specify different actions that can be taken on each piece of selected mail in the list.

Messages are displayed with the newest message at the top of the list. Urgent messages have an exclamation point next to them, and messages with an attachment are shown with an appendage icon next to the message icon.

The single most serious shortcoming of the Microsoft Mail system is that only one file may be attached to each message. This can be worked around by using one of the various compression and archiving utilities such as StuffIt or Compact Pro, but this requires additional action.

Mail is read by selecting the mail item and clicking the Read button or by double-clicking on the desired mail item in the

list. This action opens the actual message, similar to the Phone message shown in Figure 4.27.

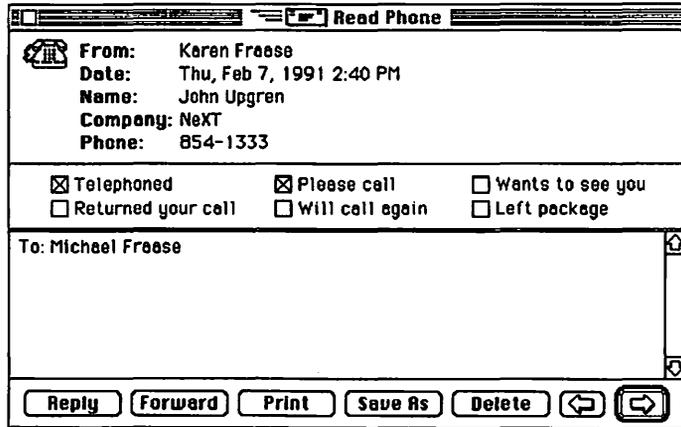


Figure 4.27 Microsoft Mail message window.

One of the most useful features of Microsoft Mail are the two navigational buttons in each message window. These buttons, similar in appearance and action to HyperCard's next and previous buttons allow you to sequentially navigate your messages. This is especially useful in workgroups that exchange a lot of electronic mail messages.

Preferences

Where QuickMail offers a wide assortment of configuration preferences (7 different windows-full), Microsoft Mail preferences are limited to a single window. As shown in Figure 4.28, the preference settings are clear and straightforward.

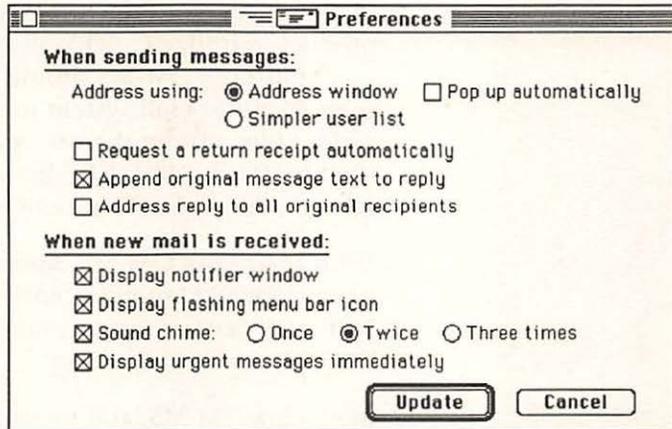


Figure 4.28 Microsoft Mail preferences.

In keeping with Microsoft Mail's highly structured approach to electronic mail, workgroup members cannot define their own groups. Nor can users create their own electronic mail forms. Forms are designed using HyperCard and installed in the mail server by the network manager.

Similarly, while bridging functionality is built into the Microsoft Mail system, no bridges are provided with the package. Workgroup members are restricted from accessing any installed bridge until the network manager has added their user account to the list of users with access to each bridge.

Evaluating Electronic Mail Systems

There are more than 330,000 Macintosh users connected to a local area network electronic mail system according to early 1990 figures from Dataquest. Of those, 180,000 or 54 percent, use QuickMail; 70,000 or 21 percent use Microsoft Mail; and the remaining 80,000 or 25 percent use something else.

During my writing cycle for this project, Berkeley, California-based StarNine Technologies Inc. was working on a product to bridge QuickMail and Microsoft Mail. The product, Mail*Link QM-MS should now be available, and allows users of either mail system to exchange messages and attached files among themselves. The only limitations with the Mail*Link QM-MS bridge are that forms are not translated and address books are not extended across both systems.

With Mail*Link QM-MS, Microsoft Mail servers appear as normal QuickMail mail centers to QuickMail users. QuickMail mail centers are accessible to Microsoft Mail users through a built-in gateway.

Mail*Link QM-MS, and products like it that are sure to follow, makes your choice of an electronic mail product less crucial.

With that in mind, it's important to choose the electronic mail system that is the most appropriate for your work-group. QuickMail and Microsoft Mail are both worthy of consideration but approach the concept of electronic mail from opposite ends of the spectrum.

Aside from the cosmetic and user-customization issues mentioned earlier, the biggest difference between the two electronic mail systems is the way they handle the message base. QuickMail stores each message as a separate file and creates multiple folders at various locations on both the server and user computers. This leads to disk fragmentation but offers the advantage of at least a chance of partial recovery in the event of a disk error. Microsoft Mail, on the other hand, creates only a single file on the mail server's disk and stores each mail message in this single database file. The advantage to this approach is lowered disk fragmentation, but if the mail file is damaged, all mail is lost.

The current version of Microsoft Mail is severely hampered in some applications by its inability to handle more than a single file enclosure. Microsoft has announced an upgrade

that the company says addresses this shortcoming. Microsoft Mail also lacks real-time conferencing capabilities; a feature that at first glance looks gimmicky but quickly proves its worth in a workgroup setting.

Microsoft Mail is more efficient in its handling of copies of the same message sent to multiple recipients. QuickMail sends—and stores—a copy for each addressee where Microsoft Mail sends only a single copy of the message.

Microsoft Mail also is more flexible in support of varied server environments. QuickMail requires that the mail server reside on a Macintosh while Microsoft offers a mail server that operates on a Digital Equipment Corp. VAX. Both QuickMail and Microsoft Mail can support IBM PC clients.

QuickMail's user customization features are significantly stronger than those found in Microsoft Mail. QuickMail allows each user to define his or her own groups; Microsoft Mail allows only the network manager to define groups. QuickMail allows workgroup members to sort messages by priority, subject, date, or sender. Microsoft Mail, on the other hand, offers message sorting by date only. QuickMail also provides the ability to unsend a message, a relatively minor feature that Microsoft Mail lacks.

Performance of both electronic mail systems is virtually identical on networks that carry a normal load of traffic. On networks with heavy traffic, however, Microsoft Mail can be significantly faster than QuickMail, especially on messages with attached files.

The pricing structures of the two products is also significantly different. Microsoft charges \$395 for the Microsoft Mail server software and \$125 for a four-user package of the client software. CE Software charges \$339.95 for a five-user package that includes the server. For larger workgroups, the pricing structures are just as disparate. A Microsoft Mail 20-user package retails for \$1,495 and the \$395 server package is still

required. Two QuickMail 10-user packages retail for \$499.95 and do not require a separate server purchase.

Because of QuickMail's included bridges, remote access, real-time conferencing, user customization features, and price advantage it is clearly the most appropriate choice for most workgroups. The exceptionally high quality of CE Software's technical support is also legendary in the industry.

Network Management

Network management for a collaborative workgroup is a set of varied tasks that are usually relegated to a network manager or administrator. Small workgroups generally don't have the luxury of a dedicated network manager, and the job usually gets passed to one of the actual workgroup members. This is not a problem until the tasks of administering and managing the network grow to such an extent as to interfere with the productivity of the acting manager. At that time it's important to either hire a full-time manager, delegate the network management tasks among the workgroup members, or seek help on a contract basis from a knowledgeable freelancer.

The network management tasks themselves are widely varied, ranging from diagnosing network problems to tweaking performance of the network, to technical support services for workgroup members. Because of the wide variety of the tasks involved with managing the network, no single program is available that is capable of performing all of the tasks. The network manager must assemble a collection of tools to aid in his or her tasks.

The International Standards Organization (ISO) has defined five categories of network management tasks:

1. *Fault management.* The process of diagnosing and remedying abnormal events on the network. Fault manage-

ment tasks are usually of a troubleshooting nature: a printer that disappears from the network or a file server that crashes, for instance.

- Useful tools for fault management tasks are *network pollers* such as Apple Computer's InterPoll, Farallon Computing's NetAtlas or CheckNet, and Technology Works' GraceLAN. A poller actively checks for devices on the network.
 - Another kind of fault management tool for network management is a *traffic monitor*. Rather than checking for specific devices, traffic monitors listen to all of the message packets carried on the network and display a graphic representation of the network's traffic. Farallon Computing's Traffic-Watch II is an example of a traffic monitor.
2. *Performance management*. The process of analyzing and enhancing the overall performance of the network. Performance management tasks are usually performed infrequently, on an as-needed basis when network performance degradation becomes noticeable.
- Traffic monitor tools such as Farallon's Traffic-Watch II can be used to identify a network device that is clogging network traffic by watching for unusually high numbers of message packets being sent from any device.
 - If the traffic monitor tool shows that none of the network devices is clogging network traffic, you can use a *packet analyzer* to diagnose the problem. Like a traffic monitor, a packet analyzer listens to all network traffic, but provides a much deeper level of information about each message packet. Packet analyzers are generally quite expensive and require a level of expertise that may require outside assistance.

3. *Configuration management.* The process of managing the interoperability of the various network devices. Software version maintenance, technical support services, and other general facilities management tasks also fall under this category.
 - The easiest—but most limiting—way of assuring compatibility between various network devices is to purchase all of your routers, gateways, and other network devices from a single manufacturer. Shiva Corporation, for example, offers its Internet Manager which allows you to configure all Shiva devices from a single software program.
 - A more appropriate solution for most workgroups—because it allows a richer diversity of equipment to be used—is to use one of the newer network management applications such as Pharos Technologies Inc.'s Status*Mac or Technology Works' GraceLAN. Both products offer extensive profiling and analysis of all network devices on the network (or internetwork) and are relatively simple to use.
4. *Security management.* The process of ensuring the security and integrity of the network, its components, and its data. This is not as great a problem as it was in the past; most of the newer routers, for example, have options that allow the manager to specify which network zones are available to which users.
 - Remote network access via modem raises an entirely different set of security issues. Any time any part of a network is accessible via outside telephone lines, security becomes an important issue because access is no longer limited to geographic proximity. Two levels of security are available for remote access. The first level involves assigning a simple password for remote access to the network. The second, more secure level of security involves

the use of *callback* features. Callback accepts a call from a remote workgroup member, hangs up the phone, and redials a preassigned telephone number for that user to re-establish the connection. As of this writing, Farallon Computing's PhoneNET Liaison is the only router that offers a callback feature. Liaison offers the additional benefits of maintaining a log of all remote access to the network, and controlling the zones accessible to a remote workgroup member.

5. *Accounting management.* The process of tracking the use of various network resources. Ideally, accounting management tools for network management would include logs of printer use, shared modem use, file server storage space usage, and the like. Unfortunately Macintosh network resource accounting management utilities are quite rare.
 - The AppleShare file server administration program, for example, provides only limited usage information. The QuickMail electronic mail system, on the other hand, provides a fairly extensive report of Email usage.

The following sections cover the best Macintosh network management tools available.

GraceLAN

The best place to start in assembling your toolbox of network management utilities is Technology Works' GraceLAN. While GraceLAN doesn't provide all of the functionality necessary for managing, administering, and maintaining a local area network, it comes closest. By allowing a single person to collect detailed information about all the devices connected to the network, GraceLAN is usually the first place to start for diagnosing network anomalies. It's also indispens-

able in configuring the network and providing technical support to other workgroup members.

GraceLAN is installed on every computer attached to the network by copying a startup document (INIT) named GraceLAN Responder into the system folder of each Macintosh and restarting. IBM PCs and compatibles are also supported by installing the GLRESP.COM file and modifying the AUTOEXEC.BAT file on each PC.

For the sake of simplicity, all figures in this section will be based on a simple two Macintosh, one LaserWriter network.

After the GraceLAN Responder has been installed on all the computers, the network manager uses the GraceLAN application to obtain various levels of information about the network and the configuration of each workstation. Launching the GraceLAN application and selecting the default zone name results in the initial GraceLAN display shown in Figure 4.29.

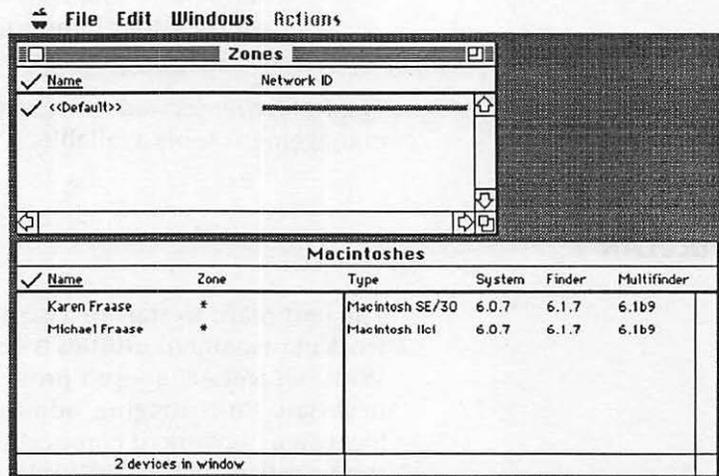


Figure 4.29 Initial GraceLAN display with the default zone selected.

All available zones on the network are listed in the Zones window and basic information for each Macintosh in each zone is shown in the Macintoshes window. The Macintoshes window provides the following at-a-glance information about each Macintosh in the currently selected zone:

- Type of Macintosh
- System, Finder, and MultiFinder version numbers
- LaserPrep version number
- Amount of installed memory
- AppleTalk version number
- MultiFinder state
- Active application
- Full user name
- Last startup date and time
- Access setting of the GraceLAN Responder
- AppleTalk Node number

This is a very useful amount of information that is available about each Macintosh that is applicable for technical support, and network diagnosis. Specific applications for GraceLAN will be explored in Chapter Five: Tutorials.

By activating any item in the Macintoshes window and selecting the Mac Details command on the Windows menu, the network manager can obtain detailed information about the selected workstation unobtrusively, and without any action on the user's part at all. This is invaluable information in many circumstances, and the unobtrusive nature of the network polling is what sets GraceLAN apart from similar utilities. The detailed information for a single workstation is shown in the following illustrations, Figure 4.30 and Figure 4.31.

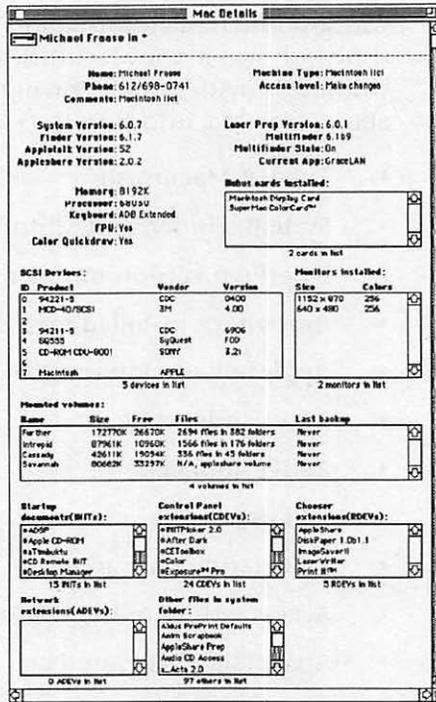


Figure 4.30 GraceLAN workstation detail window.

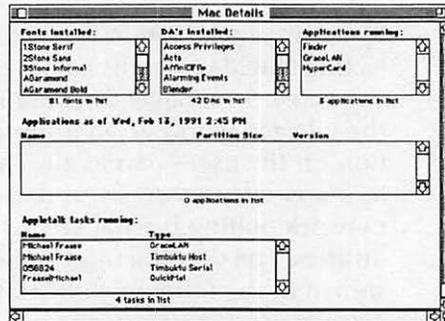


Figure 4.31 GraceLAN workstation detail window extended.

As you can see, a tremendous amount of information is provided in the GraceLAN Mac Details window. Even though the window is shown in two parts in Figure 4.30 and Figure 4.31, the window is actually one large scrolling window.

The information provided in the details window is enough to provide the network manager with a complete hardware and software inventory, including all attached peripherals, installed fonts and desk accessories, control panel devices, startup documents, and chooser extensions.

A topology of the entire network is also available to the network manager. Selecting the Topology command from the Windows menu displays a graphic overview of the entire network, as shown in Figure 4.32.

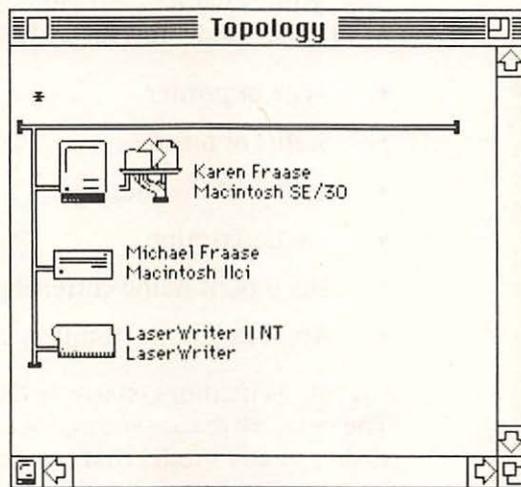


Figure 4.32 GraceLAN topology window.

A hierarchical outline display of the network topology is also available by clicking the small Macintosh icon in the lower left corner of the topology window. Additionally, dou-

ble-clicking on any Macintosh in the topology window will display the details window for that computer.

Selecting the Printers command from the Windows menu will display a window listing all available printers. An example is shown in Figure 4.33.

✓ Name	Zone	Type	Status	Source	Job
✓ LaserWriter II NT	*	LaserWriter	waiting	AppleTalk	

1 device(1 selected) in window

Figure 4.33 GraceLAN printers window.

The Printers window provides the following at-a-glance information about the currently selected printer:

- Type of printer
- Status of printer
- Source of the message
- Job description
- Document name currently being printed
- AppleTalk node number

GraceLAN includes relatively extensive analysis capabilities. The network manager can, for example, search for all Macintoshes of any model that have a certain type of NuBus card in common. Selection criteria can also be combined, resulting in a fairly complex search that returns all Macintoshes with Apple 8•24GC cards and System 6.0.7, for instance.

Invasion of privacy issues are a great concern in any networked environment. GraceLAN addresses this issue admirably by allowing each workgroup member to have complete control over how much (or how little) informa-

tion about their workspace is provided to the network manager. Four different levels of security are provided in the GraceLAN Responder startup document as shown in Figure 4.34.

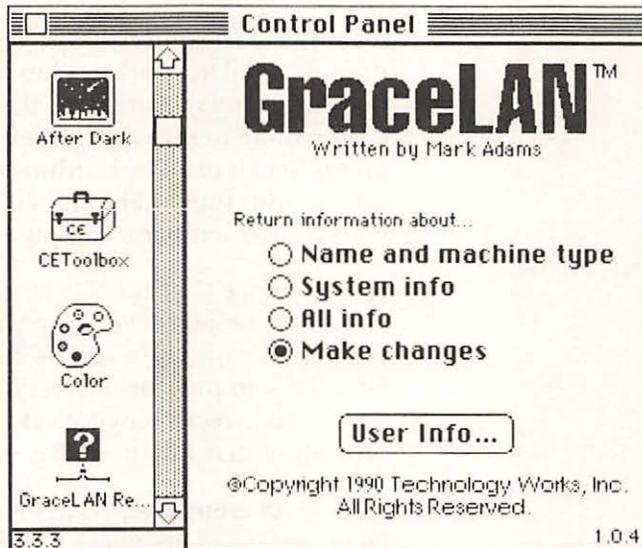


Figure 4.34 GraceLAN Responder security levels.

- *Name and machine type.* Only the type of the computer and its network name is returned to the network manager.
- *System info.* Returns the network name, machine type, and information about the contents of the System folder.
- *All info.* All information that is retrievable by GraceLAN is available to the network manager.
- *Make changes.* All information is available to the network manager. In addition, the system clock can be reset and the GraceLAN Responder can be updated.

Workgroup member documents and data files are protected from the prying eyes of GraceLAN regardless of the security setting in the GraceLAN Responder.

GraceLAN is impressive in its thoroughness, although two added functions would make the product even better.

Ideally, the GraceLAN Responder would supply the serial number of all network devices being polled. Because there is no unique serial number in the Macintosh ROM, it's likely not possible to provide this level of functionality. A comments field is provided within the GraceLAN Responder that can contain the serial number. This would require each workgroup member to supply the number, however.

Second, since GraceLAN is already capable of updating the GraceLAN Responder on any Macintosh across the network, it would seem to be a simple matter for future versions of GraceLAN to provide universal software updating capabilities. In fact, Technology Works has already begun working on adding that functionality.

Some workgroups, especially large ones, will bemoan the fact that GraceLAN doesn't support the database approach to network management. GraceLAN's creators felt that support for real-time inquiry was more important than a complete database. Both approaches require tradeoffs, to be sure, but similar products that feature a database are significantly slower than GraceLAN and cannot be used to provide real-time technical support or provide a snapshot of the condition and topology of the network.

GraceLAN is a good tool for providing technical support for workgroup members. Because it is unobtrusive, network managers and support personnel can get an in-depth profile of a user's machine without so much as asking the user to open their System folder.

GraceLAN is useful for general network management and administration as well as the following:

- Inventory of all installed hardware and software, including version information.
- Identify potential software conflicts.
- Obtain information about obsolete software versions still in use on the network.
- Synchronizing all clocks on the network.
- Create a graphic or hierarchical topology map of the network

GraceLAN pricing is \$250 for 50 users and one administrator. Additional 50 packs are available for \$195.

NetUpdater

Until GraceLAN is capable of universally updating software applications across the network, the next best thing is MDG Software's NetUpdater.

NetUpdater, from MDG Computer Services of Arlington Heights, Illinois, is a HyperCard-based application that lets a network manager update files and applications, automatically, over the network. Specifically, NetUpdater can be used to automatically perform the following network management tasks:

- Update QuickMail address books, forms, and groups
- Update printer drivers
- Update anti-virus utilities
- Update any application
- Create folders and copy any application or document into the new folder
- View a list of all workgroup members on the network

NetUpdater, used in its simplest form, requires the software to be installed on each Macintosh on the network. To automate the updating process, NetUpdater must be set as the startup application. In cases where this is inconvenient or when you don't want to burden the network load with the updating process, a separate application is provided. The NightUpdater stack carries out the update process during the night, after everyone has left for the day.

Because NetUpdater is a HyperCard stack, it is easily modified for custom applications. In fact, the NetUpdater distribution package includes NetUpdater Scriptor that lets you customize the updating process in virtually any manner you wish.

NetUpdater is installed by placing a copy of the NetUpdater stack in the System folder of each Macintosh on the network. Both NetUpdater and NetUpdater Scriptor stacks must be installed on the network manager's Macintosh. The HyperCard application and Home stack must also be installed on all Macintoshes.

NetUpdater works by comparing the dates of the source items specified for update on the file server with the date of destination items of the same name on each workstation. If the date of the source item on the server is later than the date on the workstation, NetUpdater automatically performs the update when the NetUpdater stack is launched on the workstation.

The NetUpdater stack can be launched in any of three ways:

1. NetUpdater can be set as the startup application
2. NetUpdater can be manually opened within HyperCard
3. NightUpdater can be used to launch NetUpdater at a random time over night.

When NetUpdater is launched it checks to see if an AppleShare file server volume is mounted. If a mounted server vol-

ume is not found, NetUpdater logs into the file server as a special NetUpdater user. When NetUpdater has successfully logged into the AppleShare server, it checks for any new update scripts. If new updates are available, NetUpdater performs the updates and creates a log file on the file server. The log file contains version information about System software, printer drivers, startup documents, free disk space, installed memory, and SCSI devices.

The network manager does not have to create a separate NetUpdater script for each workgroup member. NetUpdater is capable of recognizing Groups created by the network manager. NetUpdater groups can be configured in the same manner that AppleShare groups are set up, or they can be configured in a completely different manner.

There are two types of NetUpdater scripts, "Run Once" and "Run Frequently." The "Run Once" script is most useful for updating a new version of a specific software program, or other items that only need to be installed once. The "Run Frequently" script is ideal for items that change on a regular basis, such as QuickMail address books, company style sheets, site dictionaries, and the like.

NetUpdater works reliably so long as all the Macintosh clocks on the network are set accurately, but updates of large software applications can cause significant performance degradation of the network. Most workgroups will find NetUpdater useful for updating software automatically during the times when network activity is the lowest, usually when everyone has left for the day.

The most significant drawback to NetUpdater is its price. Although an evaluation package is available for \$100, the pricing structure for NetUpdater is steep. The smallest package available, with support for up to 50 users retails for \$1,000. Support for up to 100 users is \$1,700. NetUpdater is an appropriate choice for large workgroups (or corporate sites with many workgroups), if only because it's the only option available as of this writing. There are other products current-

ly under development that will perform at least as well as NetUpdater at a significantly lower price. Small workgroups would be well advised to wait.

Teleconferencing

In the 1960s, the office of Emergency Preparedness applied some of Doug Engelbart's early research in the implementation of a remotely-accessible computer-supported workgroup system for responding to national emergencies. The system itself was based on the Delphi method of developing expert consensus. In the late 1960s, this was a quite forward-thinking application of computer technology and laid the foundation for what we now commonly refer to as teleconferencing.

In the mid-1970s, Murray Turoff at the New Jersey Institute of Technology began work on a remotely-accessible collaborative system that would provide support for workgroup research.

Most observers and researchers agree that it is too soon to apply a rigid definition to teleconferencing because the technology is changing so quickly. For the purposes of this work, teleconferencing means interactive group communication through any electronic medium.

Teleconferencing holds a great deal of promise for collaborative computing. Workgroup members will be able to meet whenever the need arises without having to leave their geographic locations. This will allow more heterogeneous workgroups to be formed because physical proximity will be less and less of a factor.

As the computer and video technologies continue to merge, readily accessible teleconferencing systems will evolve from the text-oriented systems we have today to systems with full-motion video and full-spectrum audio from multiple

sources. Make no mistake, these systems already exist, it's just that the attendant high cost puts such systems out of the reach of most workgroups. Apple Computer and other companies have already demonstrated the technology for such systems that are based on readily available computer and video components.

Teleconferencing can be subdivided into two basic forms:

- *Synchronous teleconferencing.* All conference participants are present simultaneously, regardless of physical location, in real time. Telephone conference calls are the most widely used example of synchronous teleconferencing.
- *Asynchronous teleconferencing.* Conference participants check into the ongoing conference at whatever time is convenient for them. A voice mail messaging system or electronic bulletin board system (BBS) are widely used examples of asynchronous teleconferencing.

Synchronous teleconferencing facilities for the Macintosh are available only in limited, demonstration and prototype forms in various development labs. There are various applications and services that offer asynchronous teleconferencing on the Macintosh, however. The national telecommunications-based information services such as CompuServe, GEnie, and America Online, for example offer various forms of asynchronous teleconferencing. Several electronic bulletin board systems also exist for the Macintosh, and one stands out in its approach and implementation.

TeleFinder

TeleFinder, from Spider Island Software of Irvine, California, is unique in its approach to Macintosh-based telecommunications. Most Macintosh teleconferencing systems implement a command-line interface or a graphical interface laid

on top of a command line. TeleFinder attempts to emulate the Macintosh Finder completely in its interface and the underlying metaphor is that of adding telecommunications features to the basic Finder interface. Conferences are represented as folders, files are transferred by dragging them from disk to disk, and conference messages can consist of anything you can do on the Macintosh: text, graphics, animation, or sound. Provision is also made—with a command line interface—for workgroup members that are using computers other than the Macintosh.

TeleFinder's emulation of the Macintosh Finder for its interface is quite effective. Launching the TeleFinder/User application causes the main TeleFinder window to be displayed as shown in Figure 4.35.

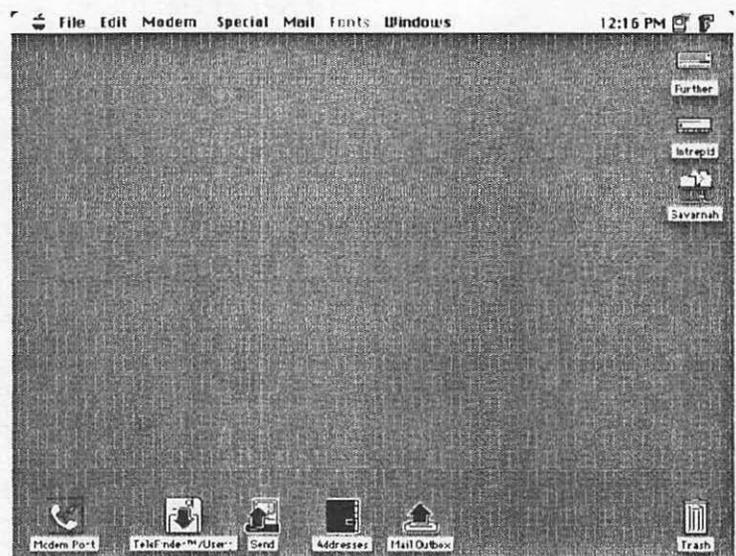


Figure 4.35 *TeleFinder main display.*

To initiate a telecommunications session and automatically log into the conference host the workgroup member double-

clicks the modem port icon shown in the lower left corner of Figure 4.35. Doing so causes TeleFinder to display its Auto Dialer window shown in Figure 4.36.

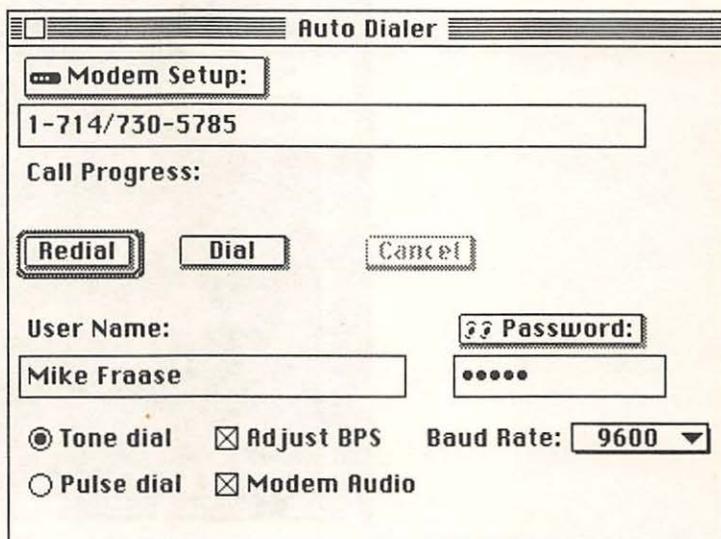


Figure 4.36 Telefinder Auto Dialer window.

The Auto Dialer settings are usually pre-configured by the network manager before the TeleFinder/User application is given to the workgroup member. In ordinary operation all the workgroup member needs to do to connect to the host is to click the Redial or Dial buttons to initiate the telephone call. Options are provided for changing the modem configuration, the telephone number of the conference host, and the user name and password.

The log in sequence is performed automatically. When the workgroup member is connected to the remote TeleFinder host, the screen is updated to include additional icons for disks, conferences, and the workgroup member's mailbox

on the remote host. The updated display is shown in Figure 4.37.

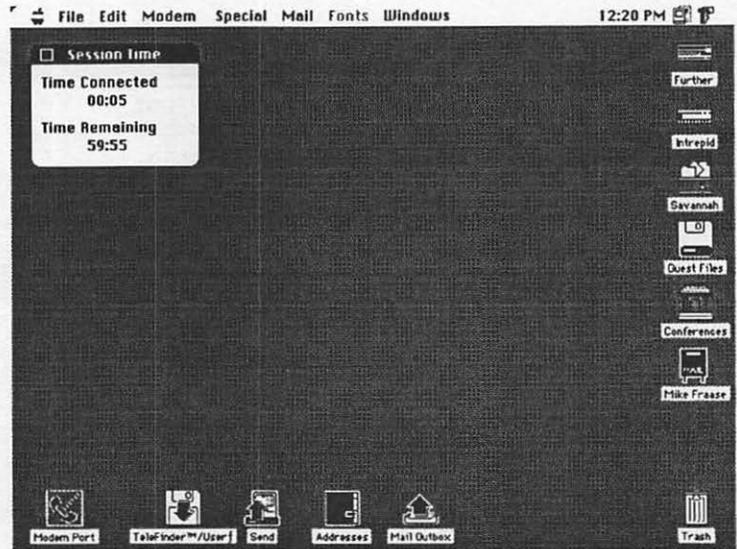


Figure 4.37 *TeleFinder user display upon connection to the host.*

Once connected to the host, TeleFinder works in very much the same way as the native Macintosh Finder. The disks and conferences on the host can be opened by double-clicking on them, and files can be transferred by dragging them from disk to disk. An elapsed connection time counter is also available.

Double-clicking on the Guest Files disk opens that folder on the host. The host disk may contain multiple folders, and even multiple disks can be available, each represented by its own icon within the TeleFinder display. Opening the Guest Files disk results in a display like the one shown in Figure 4.38.

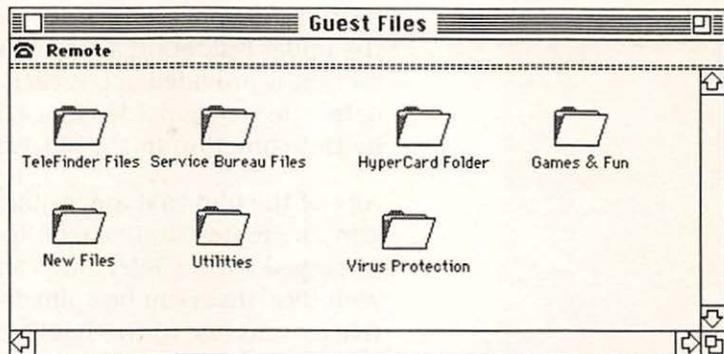


Figure 4.38 *TeleFinder remote disk window.*

The window, with minor exceptions looks very similar to a normal Finder window. The exceptions are that Folders that are available to the remote user are displayed with a partially open folder icon. Also, a small telephone icon and the word Remote appear in the upper left corner of the window to indicate that the folder resides on a remote volume.

The folders on the remote host's volume work pretty much as you would expect. Double-clicking on one of the folder's opens it in a new window, such as the one for the Service Bureau Files folder shown in Figure 4.39.

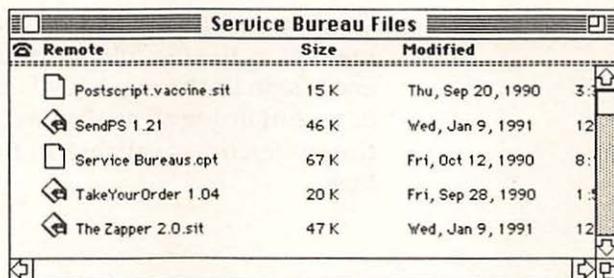


Figure 4.39 *TeleFinder remote folder window.*

Again, the small telephone icon and the word Remote appear in the upper left corner of the window to indicate that the folder resides on a remote volume. Additionally, information is provided about each file's size and modification date. The view provided for each folder is similar to the View by Date function in the Macintosh Finder.

Any of the files that are visible can be downloaded by dragging the representative icon to one of your local disk icons displayed within TeleFinder. Similarly, any file on one of your local disks can be uploaded by dragging the representative icon to one of the TeleFinder host disks.

TeleFinder supports Zmodem file transfers with complete failure recovery. The Zmodem file transfer protocol is capable of resuming a failed file transfer from the point at which the original transfer failed. TeleFinder adds to the transfer recovery capabilities of Zmodem by maintaining a separate transfer process log for each user. If a file transfer fails, it will automatically resume at the point of failure the next time you log into the remote TeleFinder host, regardless of whether it is minutes, hours, or even days later.

Any number of conferencing sections can be set up on the remote TeleFinder host by creating new folders for each section. In turn, these conference sections can contain subconferences that are set up as subfolders within the predefined conference folders.

Conferences are opened on the TeleFinder host in the same manner as the file folders: double-clicking on the Conferences icon in the main TeleFinder/User display opens a window containing all available conferences. Figure 4.40 shows the conferences available on the example remote TeleFinder host.

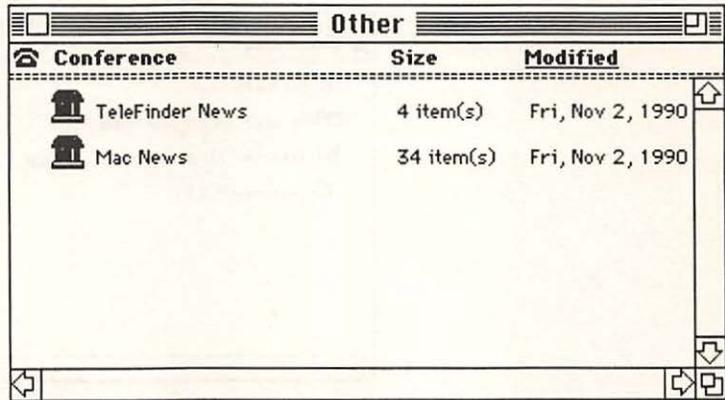


Figure 4.40 TeleFinder conferences window.

Once again, the TeleFinder conference window looks relatively similar to an ordinary Finder window. The exception is the small telephone icon and the word Conference appearing in the upper left corner of the window to indicate that the folder resides on a remote volume.

The number of conference items currently available within each conference is listed as well as the last modification date for the conference. The conference itself may contain any type of material that can be created on the Macintosh, including text, graphics, animation, or sound files.

To enter a conference, the workgroup member double-clicks on the appropriate icon. In the example shown in Figure 4.40, two conferences are available: TeleFinder News and Mac News. To enter the TeleFinder News conference, for example, the workgroup member would double-click on the TeleFinder News icon, resulting in a listing of materials available in the conference shown in Figure 4.41.

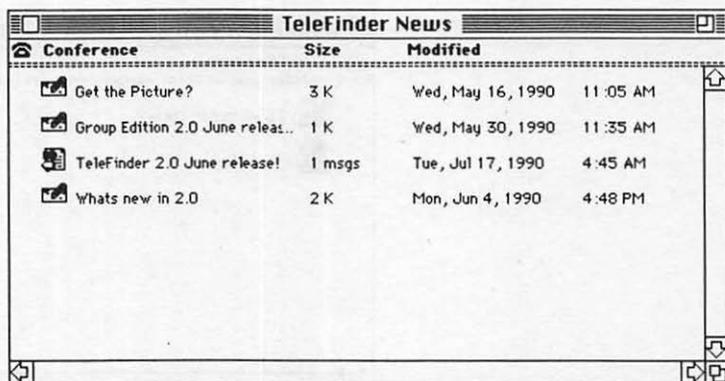


Figure 4.41 TeleFinder conference item list.

Any conference item that can be viewed online can be opened by double-clicking on it. A text message, for example, is displayed like the one shown in Figure 4.42.

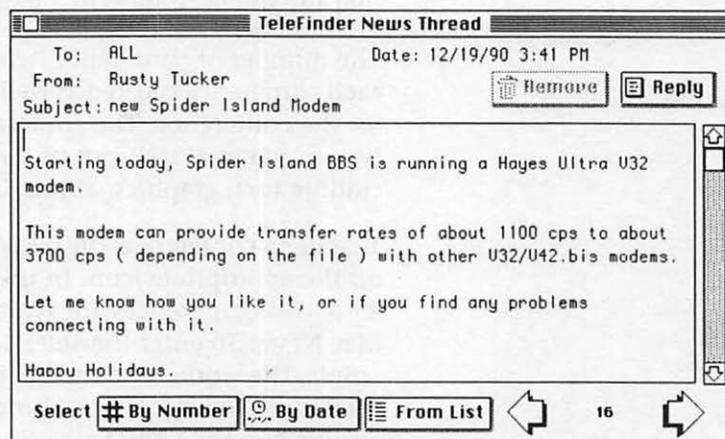


Figure 4.42 TeleFinder text conference item.

Buttons are provided within the conference item window allowing the workgroup member to remove the current item, reply to the current topic, select an item by number, by date, or from an abbreviated list of conference items. Navigational buttons are also provided for moving to the next or previous conference items.

As simple as TeleFinder is to use for remote members, it is somewhat surprising that the TeleFinder host application is just as easy to configure and maintain. The TeleFinder host application, called Group Edition Host, automatically opens to a default state of waiting for an incoming call when launched.

Selecting the Configuration... command from the File menu brings up a window that allows the network manager to specify the conferences and file areas available to the workgroup members. A sample configuration is shown in Figure 4.43.

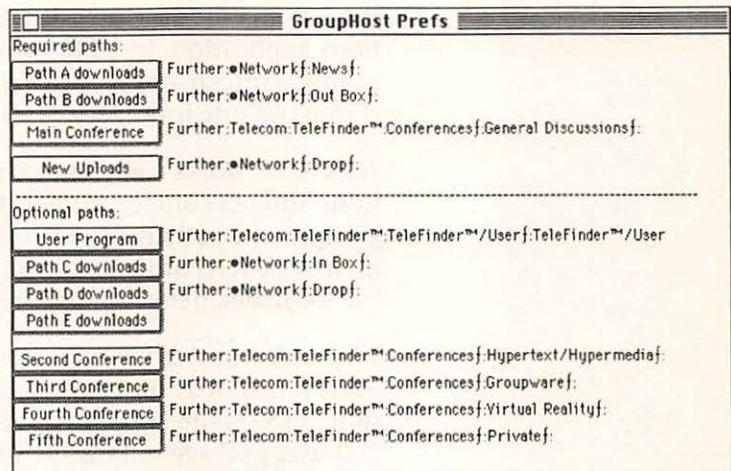


Figure 4.43 *TeleFinder Group Edition Host configuration window.*

The top four items in the configuration window are required for the TeleFinder host to operate properly; the items below the dividing line are optional.

Any configuration item can be modified by clicking on the button associated with the path for the conference or file section. Clicking on any of the configuration buttons brings up a standard Macintosh file dialog box that allows the network manager to specify a new path for the selected conference or file area.

The only other configuration items within the TeleFinder Group Edition Host are modem-specific settings such as the initialization string, baud rate, handshaking method, and the like. TeleFinder comes with preconfigured drivers for most standard modems, and the product works well with the newer high speed modems.

Because TeleFinder requires specialized software on both ends, an option is available that allows workgroup members to download the necessary TeleFinder/User application from the Group Edition Host with any standard telecommunications application. In addition, complete access is afforded non-Macintosh workgroup members through a DOS-like command line interface.

TeleFinder works well, and helps demystify telecommunications and teleconferencing. It is appropriate for most workgroups, especially those with geographically dispersed members. TeleFinder would be even better if it worked over the local area network as well as telephone lines.

TeleFinder is priced by the number of callers that will be supported. A 20-user package is available for \$375; a 50-user version is \$575; the 100-user configuration costs \$895; and a 200-user package sells for \$1,250. The 20-user version is capable of supporting most workgroups, and the 100- and 200-user versions are most appropriate for enterprise-wide teleconferencing.

Remote Network Access

Another approach to providing asynchronous teleconferencing services for a collaborative workgroup is to allow remote access to the local area network itself.

The disadvantages of allowing remote access to the network include decreased security for the network itself, and reduced performance on the remote link. Most Macintosh-based network services were not designed for remote access.

The advantages of remote network access is that remote workgroup members can be afforded access to the entire set of services provided on the local area network, including printer service, file service, and electronic mail service.

Teleconferencing can be implemented by creating the appropriate folders on the file server and defining the appropriate electronic mail accounts.

PhoneNET Liaison

Farallon Computing's PhoneNET Liaison is the best way of providing remote access to the local area network. It offers the additional benefit of providing router services for the network.

Liaison is a fully automatic software router that can be used to connect up to five AppleTalk networks of various types, including LocalTalk, EtherTalk, and TokenTalk. Liaison can also be used to provide half- or full-bridge functions between two (or more) geographically dispersed internetworks.

Liaison runs in the background of any Macintosh connected to the network and uses the resources of that computer to provide routing services for the entire internetwork. Liaison can be used to provide Ethernet users with access to a Lo-

calTalk printer such as any of the Apple LaserWriters. Conversely, it can be used to allow LocalTalk users access to an EtherTalk file server. Liaison can also be installed on a dedicated AppleShare file server Macintosh and will run in the background on the file server.

Liaison is installed by dragging the Liaison file into the System folder of the Macintosh designated to provide router services. Liaison will automatically configure itself to run on either AppleTalk Phase 1 or AppleTalk Phase 2 (or a mixed network comprised of both).

Liaison can be used to link various networks within a large enterprise. The networks can be divided by organizational department, by physical floors in the building, or by virtually any other configuration. The network manager can organize the various network areas into zones, allowing more efficient use of network services and resources. Each zone will operate faster because it will no longer have to deal with traffic from other zones.

If a modem is attached to the Macintosh running the Liaison software, it is possible to bridge the local area network with a remote network that is also running Liaison. Similarly, Liaison can be configured to allow incoming calls from remote networks or individual workgroup members. A remote workgroup member, logging into the network via Liaison and a modem appears as any other user and has access to any available network services such as printers, electronic mail, and file servers. Liaison can be configured to restrict remote access to only the Macintosh running Liaison, the zone containing the Mac running Liaison, or the entire internet.

Liaison will be most useful for large organizations with several different types of local area networks. Liaison's routing capabilities are excellent for connecting disparate networks. Smaller departments or workgroups with remote members may find Liaison useful for remote access to the local area network.

Farallon's PhoneNET Liaison retails for \$295, and allows half-bridge connections for any number of remote users. Full-bridge connections require two copies of Liaison.

Virtual Meeting Facilitation

Most of us perceive meetings as a mixture of boredom and frustration, where nothing much gets done. We want to be productive, but the typical meeting has come to be perceived as counterproductive. Most meetings are indeed counterproductive because they tend to take place in counterproductive environments.

Wouldn't it be wonderful if you could participate in meetings from your own office or workspace? It's possible to do just that with existing technology. Most meetings do not need to take place in real-time. For those situations, the teleconferencing tools discussed in the previous section can provide asynchronous meeting support.

There are situations, however, when face-to-face meetings must take place, and they have to happen in real-time. The groupware tools and techniques discussed in this book can help aid those situations.

Many meetings that take place are actually a series of one-on-one conversations. Most of these types of meetings revolve around one person showing another individual or group of workgroup members how to do something, reporting on the progress of a project, or brainstorming. With available technology, there's no reason why you should have to leave your office or workspace to participate in these real-time meetings.

Virtual meetings can take place in real-time, with any number of participants. Best of all, the meeting participants do not have to be at the same physical location. The meeting can be said to be "virtual" because there is no physical meet-

ing place. The meeting takes place in the interactions between the workgroup members at their Macintoshes. The following two sections give an overview of Farallon Computing's screen-sharing software for the Macintosh. In Chapter Five: Tutorials, these tools will be used to illustrate various implementations of a virtual meeting.

Timbuktu

Farallon Computing Inc.'s Timbuktu is the leading screen-sharing utility for the Macintosh. It allows workgroup members to view and operate each other's Macintosh over the network. A guest connected to a Timbuktu host sees on the local screen exactly what is displayed on the remote host screen. In observation mode, the guest can watch the screen activity in real-time. In control mode, the guest and the host can work together on the same document using their own mouse and keyboard.

Timbuktu has three basic modes of operation:

- *Observe.* A guest can connect to a Timbuktu host and observe all screen activity in real-time.
- *Control.* A guest can connect to a Timbuktu host and work collaboratively with the host. The guest can also control the host remotely.
- *Broadcast.* Screen activity on one Macintosh can be broadcast over the network and displayed on every other Macintosh in real-time.

Timbuktu is installed as a startup document (INIT) and associated files and a desk accessory on each Macintosh on the network. Once installed, Timbuktu is activated from the Desk Accessory menu.

Selecting the Timbuktu desk accessory causes a list of available Timbuktu hosts to be displayed as shown in Figure 4.44.

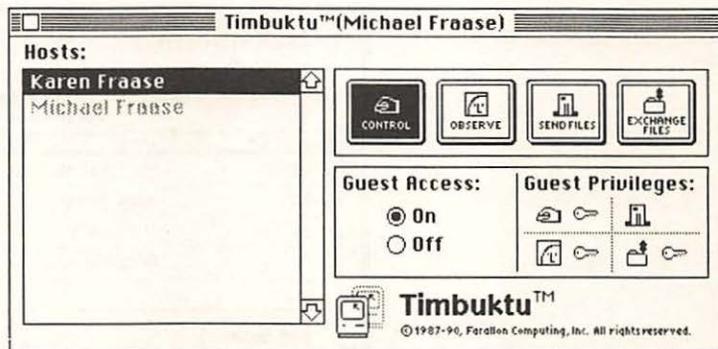


Figure 4.44 Timbuktu available hosts.

The available hosts are listed in the left panel of the dialog box. A host is selected by clicking on it. Across the top of the right panel are four buttons that determine what action to take with the selected host:

- *Control*. Obtain control of the remote Timbuktu host.
- *Observe*. Observe the screen activity of the remote Timbuktu host.
- *Send files*. Send files to the drop box on the remote Timbuktu host.
- *Exchange files*. Send files to or retrieve files from any folder on the remote Timbuktu host.

Each of the available guest privileges are defined on the host Macintosh with the Set Guest Privileges... command. Guest privileges can be defined for each workgroup member, although it is usually easier to define a group setting for various subgroups within the larger workgroup. Figure 4.45 shows a sample guest privileges setting for a small workgroup.

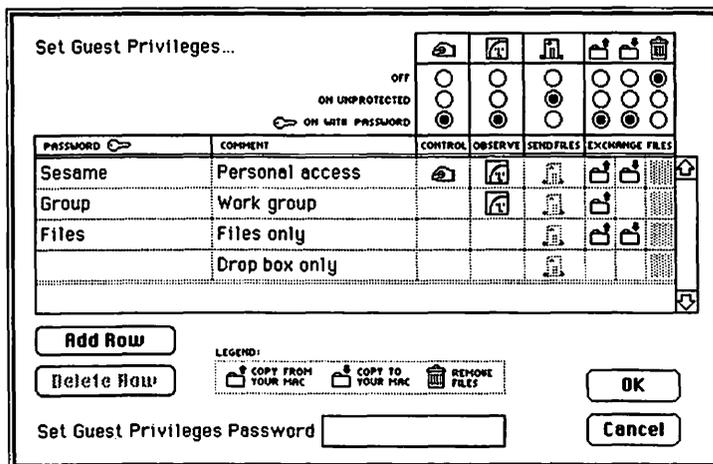


Figure 4.45 *Timbuktu guest privileges for a small workgroup.*

As you can see a different password has been assigned for three different subgroups in addition to a personal access setting. Workgroup members would be able to observe, but not control, this Macintosh unless they have the personal access password. Larger workgroups will likely need a larger set of guest privileges.

The settings shown in Figure 4.45 allow each user to set up a drop box with no password requirement. This allows anyone to send files to this Macintosh without any password. To exchange files the guest would have to know the Files password.

If all workgroup members have defined a personal access password for the control mode of Timbuktu, each workgroup member could control access to his or her Macintosh. Assuming that the personal access password is available to workgroup members that need to actively collaborate in real-time, the access method is straightforward.

To access a remote Timbuktu host you select the Macintosh you want to collaborate with by selecting it from the list of available hosts, and clicking on the appropriate action button shown in Figure 4.44. This action would result in a dialog prompting you for the password associated with the action you have selected. When you have supplied the correct password, the Timbuktu host's screen is displayed in a window on your monitor. An example of accessing a Timbuktu host in Control mode is shown in Figure 4.46.

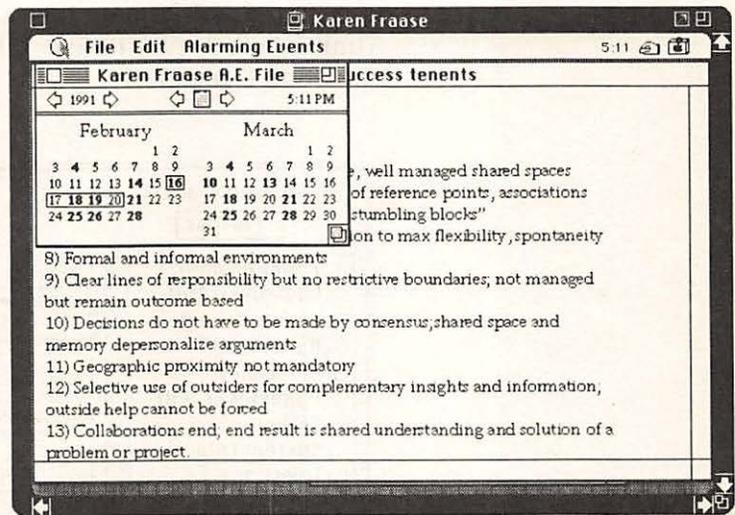


Figure 4.46 Timbuktu remote host screen in Control mode.

The Timbuktu host is notified that you have connected in Control mode by displaying the Timbuktu Control icon in the upper right portion of the menu bar (in this case, between the clock display and the MultiFinder application icon).

From this point on you can work collaboratively on a shared document on the Timbuktu host. In the example shown, an appointment calendar and a word processing document are

open on the Timbuktu host. Both the host and the guest can work interactively on the host Macintosh, although any action by the host overrides any guest action. In other words, the host retains control of his or her Macintosh and can disconnect the guest at any time with a simple menu command.

Timbuktu also offers exceptional file transfer capabilities that are significantly faster than dragging files to and from a file server. File transfers can take place in the background, allowing you to perform other tasks while the files are being transferred. Clicking on the exchange files button causes Timbuktu to display its file exchange dialog box like the one shown in Figure 4.47.

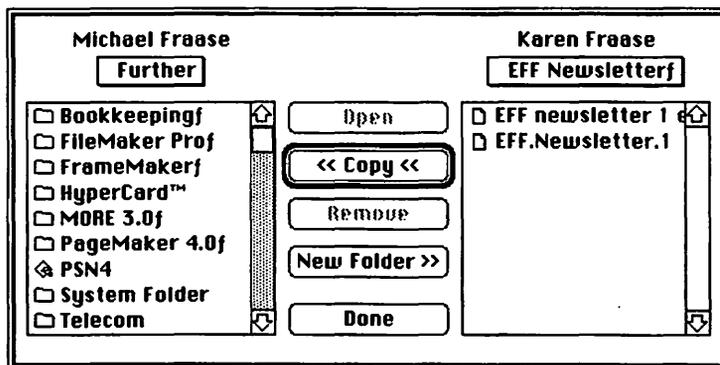


Figure 4.47 Timbuktu file exchange dialog box.

The latest version of Timbuktu, version 4.0, adds new features, including full color support between color Macintosh models and the ability to exchange Clipboard contents with other members of the workgroup during screen-sharing sessions.

Timbuktu is a powerful groupware addition to any workgroup. It retails for \$149 per computer with a 30-pack available for \$1,995.

Timbuktu/Remote

Timbuktu/Remote, also from Farallon Computing Inc. provides all the functionality of Timbuktu to remote workgroup members through a modem connection. It allows workgroup members to view and operate each other's Macintosh over standard telephone lines. A guest connected to a Timbuktu host sees on the local screen exactly what is displayed on the remote host screen. In observation mode, the guest can watch the screen activity in real-time. In control mode, the guest and the host can work together on the same document using their own mouse and keyboard.

Because modem connections are not as fast as network connections, the performance of Timbuktu/Remote is dependent upon the quality of the telephone connection and modem speed.

New features of the remote screen sharing program include:

- Support for CCL (Connection Command Language) scripting for support of more communications devices including PBXs, ISDN links, and non-Hayes compatible modems.
- Background file transfer capabilities. File transfers now take place in the background and guests can transfer files to a designated folder on the host Macintosh.
- Enhanced security features including a new password table that is used to define multiple passwords and access levels for different guests and a call-back feature. A complete log of guest activities is now available as well.
- Simultaneous viewing of both remote and local screens with the remote screen becoming a window on the local screen. The previous version took over the entire local screen.

Timbuktu/Remote is installed as a startup document (INIT) and associated files and a desk accessory. Once installed,

Timbuktu/Remote is activated from the Desk Accessory menu.

Selecting the Timbuktu/Remote desk accessory causes a list of pre-defined Timbuktu/Remote hosts to be displayed in the telephone directory shown in Figure 4.48.

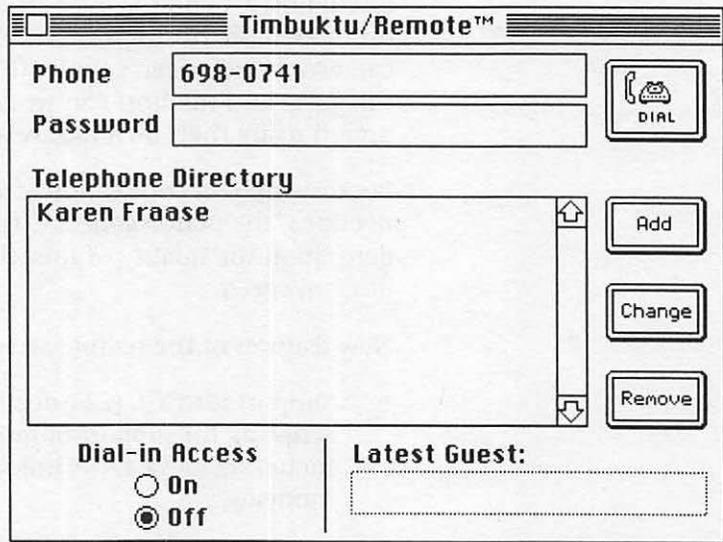


Figure 4.48 Timbuktu/Remote telephone directory.

This initial Timbuktu/Remote dialog box can also be used to turn on or off dial-in access to your Macintosh. A field also displays the last guest that logged into your Macintosh with Timbuktu/Remote of dial-in access is enabled.

To establish a Timbuktu/Remote connection with a remote workgroup member select the appropriate address book entry and click the Dial button. This initiates a dialing procedure and the automated login procedure for the Timbuktu/Remote host.

If you elect to enable dial-in access to your Timbuktu/Remote, a set of guest privileges can be defined. The guest privileges options are similar to the privilege settings for Timbuktu shown in Figure 4.45, and are set by selecting the Set Guest Privileges command from the Timbuktu/Remote menu. Guest privileges can be defined for each workgroup member, although it is usually easier to define a group setting for various subgroups within the larger workgroup.

Once the Timbuktu/Remote connection is established, operation is identical to Timbuktu used over the network. The sole exception is that screen updating, mouse actions, and file transfers take significantly longer, even with high-speed modems, because of the nature of the lower speeds available for telephone connections.

Timbuktu/Remote is a powerful utility for workgroups with remote members that need to collaborate on the same document in real-time (or as close to real-time as modem communications allow). Timbuktu/Remote is priced at \$195 per computer. Farallon also offers the Timbuktu/Remote Access Pack for \$1,295 that includes the Timbuktu/Remote software, a Remote/V.32 (9600 baud) modem, and a specialized modem cable that allows a Macintosh II-class machine to be started up automatically with an incoming telephone call.

Group Authoring & Editing

Group authoring and editing applications must overcome the underlying problem of workgroup members not sharing meanings for key concepts. Common meanings must be defined and agreed upon. This can be done in five stages of the authoring process.

- *Prewriting.* This is the stage where all the workgroup members agree on the scope and overarching goal of the project.

- *Information gathering.* Assignments are made for various content responsibilities. This helps to assure completeness with minimal overlap.
- *Organization analysis.* Quality control issues are identified at this stage and it quickly becomes clear what elements are missing and what needs to be done about it.
- *Sequence analysis.* The information gathered and analyzed is chunked into appropriate levels of granularity. This helps the sequencing issues to appear more clearly, allowing them to be dealt with more easily.
- *Presentation analysis.* Various presentation formats are experimented with at this stage. A template or series of templates is arrived at and circulated within the workgroup.

Group authoring and editing applications will be among the most important software applications forthcoming in the next few years.

MarkUp

At the time of this writing, MarkUp, from Mainstay Software of Agoura Hills, California was the only group authoring and editing software available for the Macintosh.

MarkUp uses a document center stored on a file server to provide a shared workspace for workgroup members. The MarkUp document creator, a specialized printing driver, prints a graphic representation of a file to disk rather than to the printer. This file is placed within the shared MarkUp document center for review. Associate workgroup members use the MarkUp Application to mark-up, annotate, and edit files stored within the document center.

Additionally, any document can be printed to a stand-alone reader format that can be edited, marked-up, and annotated with the MiniMarkUp application. This is useful for remote

workgroup members who do not have access to the shared document center.

Any workgroup member who has installed MarkUp can print to the shared document center by selecting the MarkUp Document Creator in the Chooser and continuing the basic printing procedure.

Most applications treat the MarkUp Document Creator as any other printing device. Issuing the Print command from within any application brings up the MarkUp Document Creator print specification dialog box shown in Figure 4.49.

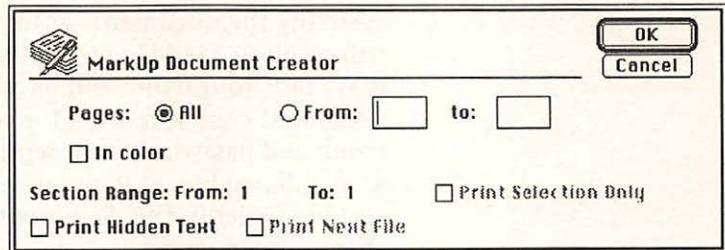


Figure 4.49 MarkUp print specification dialog box.

The MarkUp Document Creator allows you to specify the range of pages to print—or a selection—and whether or not to include color information in the document. Application-specific features, such as Microsoft Word's Print Hidden Text item shown in Figure 4.49 are also included.

When you click the OK button to initiate the printing process, MarkUp will prompt you to specify a document center to print the document to with the dialog box shown in Figure 4.50.

An option is also provided to print the document to a stand-alone MarkUp file that can be read with the MiniMarkUp application.

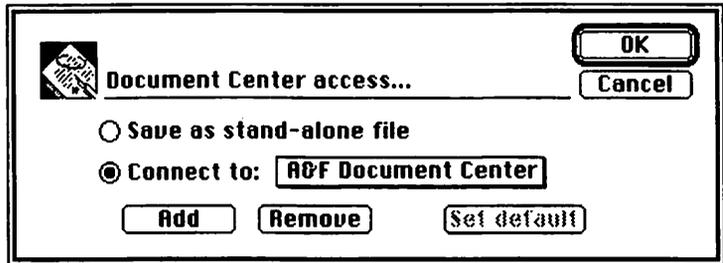


Figure 4.50 *MarkUp document center specification dialog box.*

Selecting the document center and clicking the OK button will result in MarkUp prompting you for your Name and Password. Your name and password must be known to the document center in order for you to access it. When the name and password are accepted, MarkUp will present you with a dialog box that allows you to add the document to a specific folder in the document center. This dialog box is shown in Figure 4.51.

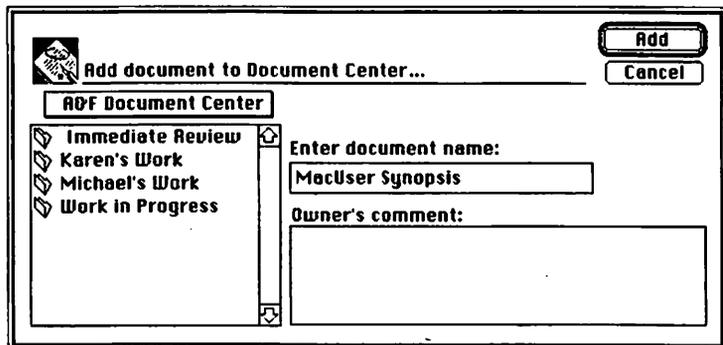


Figure 4.51 *MarkUp document center add document dialog box.*

The folder representations in the left panel of the dialog box operate like the Finder. To open one of the folders, double click it. A comment field is also provided for any comments you may wish to attach to the document.

When you click the Add button to place the document in the shared document center, MarkUp presents you with a dialog box that allows you to set access privileges and assign the document to specific reviewers. A sample MarkUp access privileges dialog box is shown in Figure 4.52.

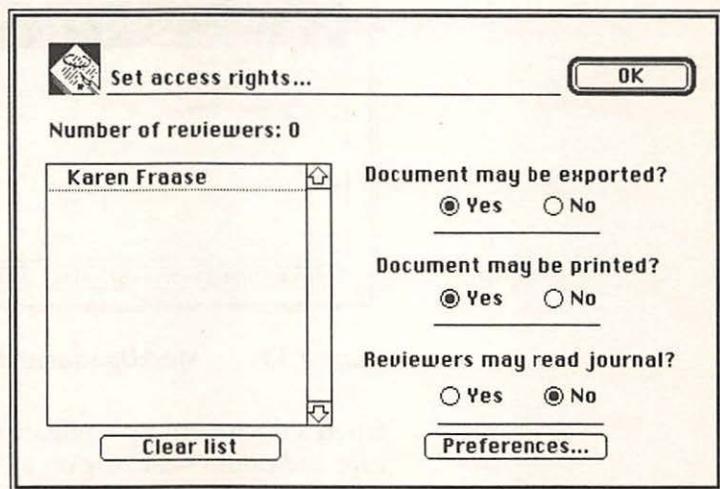


Figure 4.52 MarkUp document center access privileges dialog box.

Options are provided that allow you to restrict the reviewers from exporting the document from the document center, printing the document, or reading the review journal attached to the document.

Clicking the OK button prints the document to the shared document center where it will be accessible to the reviewers specified in the access privileges dialog box in Figure 4.52.

To access a shared document center, launch the MarkUp application. It can be set to automatically log you into a specific document center, or you may be prompted for your name and password. Once you're logged into the shared document center, MarkUp will display a dialog box similar to the one shown in Figure 4.53.

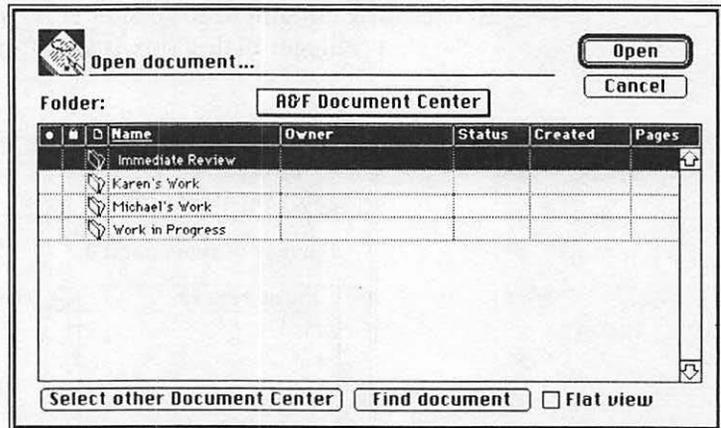


Figure 4.53 MarkUp shared document center.

Open a document by navigating through the folder structure and double-clicking on a document to open it. Only documents that are accessible to you will be displayed. Opening a document will display it within a new window.

Documents are marked up and annotated within MarkUp by adding a new overlay to the document. This is done by selecting the New Overlay command from the File menu. Doing so will add the new overlay with your name attached to it. An example of a new overlay is shown in Figure 4.54.

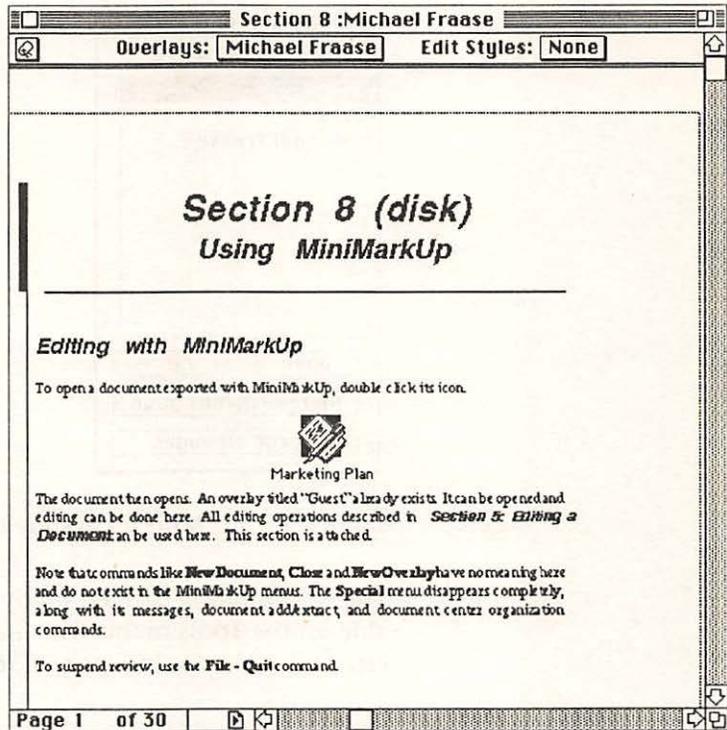


Figure 4.54 MarkUp document with a new overlay.

The document appears just as it will on paper, with dotted lines for page borders and all styling information intact. A menu for available overlays is available as is a menu for available edit styles in the upper portion of the window. Page location is displayed in the lower portion of the window.

When you create a new overlay, your name will be added to the reviewers list in the Overlays window, as shown in Figure 4.55.

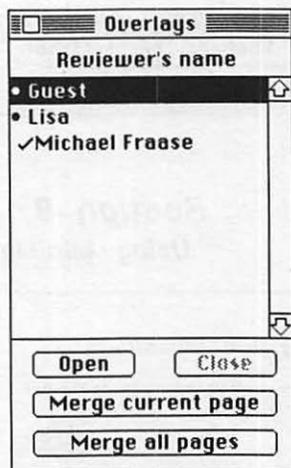


Figure 4.55 *MarkUp overlay window.*

The document is marked up by using a series of tools available on the Tools menu. This menu is actually a tear-off palette and is shown in Figure 4.56.



Figure 4.56 *MarkUp tools palette.*

The various tools are used to add different kinds of textual or graphical annotative comments to the document. Tools are available for the following:

- *Pointer tool.* Selects objects or text blocks.
- *Color palette.* Modifies the color of a selected object or text block.
- *Text tool.* Adds textual annotation to the document. The text tool can also be used in conjunction with some of the other tools.
- *Pop-up note tool.* Adds an icon representing an external note. The external note is part of the overlay and appears in its own window.
- *Line tool.* Adds straight lines, arrows, or wavy strike-out lines to the document.
- *Oval tool.* Adds ovals used to encircle portions of the document for emphasis.
- *Empty and filled rectangle tools.* Adds empty or filled rectangles to the document. The filled rectangle tool can be used as a sort of pseudo-highlighter.
- *Empty and filled polygon tools.* Adds empty or filled rectangles to the document.

Any annotation that you make to the document is stored within your overlay. The underlying document retains its original integrity, and other reviewers can add their own overlays to the document. At the end of the reviewing cycle, the original author of the document can view any of the available overlays by selecting them from the Overlays menu.

A partially marked up and annotated document, complete with a pop-up note is shown in Figure 4.57.

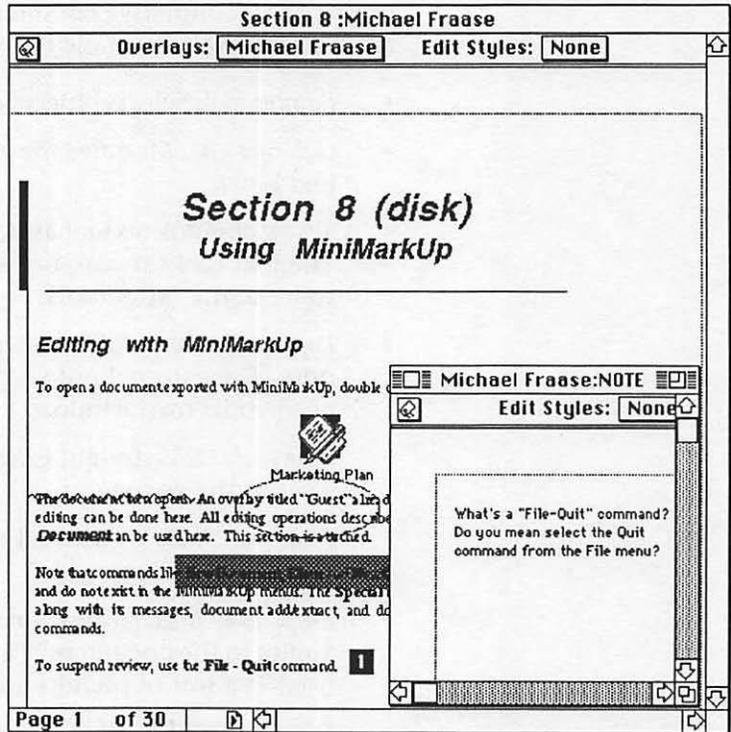


Figure 4.57 Partially annotated Markup document.

MarkUP would be more useful if it provided a richer tool set for annotating documents and if it didn't require its own shared document server. The server aspects of MarkUP force workgroup members to learn how to navigate yet another type of server mechanism, one that provides only limited usefulness.

Most workgroups will find MarkUP to be beneficial, however, as it will help cut time from the document review and revision cycle. Also, MarkUP's potential to help a workgroup conserve paper is not to be taken lightly.

MarkUp retails for \$245 per user. A two-pack is available for \$495, and a five-pack costs \$995.

Document Management

Closely related to the group authoring and editing process is the problem of document management. Where shared documents are stored, how the file server is organized, how access to various documents is managed, and how document versions are tracked are all document management issues.

Document management utilities are few and far between on the Macintosh, perhaps because the Finder provides almost adequate document management facilities for individual users. When the computing resources of a workgroup are combined, however, document management issues become critical.

Network managers can design and maintain how and where shared documents are stored on the network. It takes time, and most managers would welcome a utility that automates the task, but it can be done manually.

The versioning question is one, however, that has not been addressed. The only versioning facility built-in to the Macintosh operating system is the View by Date command in the Finder.

Several document management and version tracking applications are available for the Macintosh. Unfortunately, none of them are worthy of mention here; they are either kludges to the operating system or complex database structures.

This is certain to be a hotbed of activity in the Macintosh development community in the next few years and several products are currently under development that promise to address this market segment.

Until then, most workgroups will be well served by adopting a version numbering system to their document naming conventions. An original version of a quarterly report, for example, could bear the name Q3 Report v1.0. A minor update to the material could be named Q3 Report v1.0.1. A relatively significant update could be named Q3 Report v1.1, and so on. A complete rewrite could bear the name Q3 Report v2.0. This allows each workgroup member to be sure that he or she is working on the current version of any shared or collaborative document.

Chapter Five

Tutorials

The previous chapter, Chapter Four: Groupware Implementations, covered assessing various collaborative computing tools. This chapter provides step-by-step tutorials on actually implementing and using groupware tools and techniques in a typical workgroup environment.

Throughout this book, I have assumed that your workgroup and its needs are fairly typical. The typical workgroup consists of five to seven members with the potential for interconnection with other workgroups within the enterprise in the future. For the sake of consistency I'll assume that there are five people within the workgroup presented here and that the workgroup expects to add two more members within the next 12 months.

This workgroup's needs are typical of most workgroups: file service, printer service, electronic mail service, and screen-sharing capabilities. It would be nice to provide room for ancillary networking products, such as meeting scheduling, when they become available. Each of the five members within the workgroup have been using a Macintosh for at least three months. The workgroup's computing resources consist of the following:

- 1 Macintosh SE/30
- 2 Macintosh IIc's
- 1 Macintosh IIx
- 1 IBM PS/2
- 1 LaserWriter IINT

Additionally, budget resources are available for adding two more workgroup members within the next 12 months, one of whom will be using a Macintosh IIci; the other a NeXT-cube.

The hypothetical workgroup presented here is a documentation department within a larger enterprise, although the issues covered and the tutorials provided are also applicable to an independent firm that is comprised of less than 10 individuals.

Planning and Installing the Network

Before any sort of collaborative computing can take place, the computers must be linked together via a local area network. Before the computers can be networked, the network must be planned—regardless of how small it is. And before the network can be planned, *you must get permission from workgroup members to alter their personal workspaces*. This is critical and will prevent future misunderstandings and conflicts. Even if you are working in a corporate environment it is imperative that you get permission from all workgroup members before altering their workspace. The planning and installation phase of establishing the workgroup's basic network services can be broken down into five sequential tasks shown in Figure 5.1.

Each of these tasks will be explained in the following sections. The entire process of establishing and installing the network and basic services for all workgroup members will

take about six days. The time factors are generous; an experienced network manager can establish basic network services for a small workgroup over a weekend.

Task	Tools Needed	Time
Plan network	Documentation Site plan Network chart	2 days
Run cable	Telephone cable PhoneNET connectors PhoneNET terminators	1 day
Establish file and printer service	PSN software or System 7.0	1 day
Establish electronic mail service	QuickMail	1 day
Test network	GraceLAN	1 day

Figure 5.1 Steps for planning and installing basic network services.

Plan Network

Begin planning the network by inventorying the resources available and determining the type of network topology and media best suited for the workgroup. Since the workgroup used in this example is comprised of five members, with average storage needs, Farallon's PhoneNET in either a bus or star topology is adequate.

Next, attempt to locate a building wiring specification and blueprint for your office space. This is an important step and

can save a lot of work in the next step. Look at these materials to determine the best way to connect all workgroup members to the network in the least obtrusive manner possible.

Because PhoneNET runs over ordinary telephone wire, there's a fair chance that your cabling job may already be done for you. Many telephone systems are comprised of four wires but only use two to provide telephone service. PhoneNET is "smart" enough to utilize the other two wires for LocalTalk network service, if they're available. If this is the case in your office, connecting each node to the network is as simple as plugging the PhoneNET cable into a duplex adapter on the telephone jack in each office. These duplex adapters are available at any electronics or telephone store.

If your office uses all available telephone wiring, you need to evaluate your wiring options at this point. It may be easiest to call in a wiring specialist.

A good approach in planning your network is to sketch out a rough representation of your office layout similar to the one shown in Figure 5.2. It doesn't have to be fancy, it just has to show you the network layout and allow you to visualize your network. You can use paper and pencil, whiteboard, or a drawing program on your Macintosh to sketch out this representation.

Make note of wiring possibilities if you can't use the existing telephone wiring. Network cabling can be hidden under suspended ceilings, for example. As a last resort think about where to drill holes in the office walls.

Finally, make note of the amount of cable and number of PhoneNET connectors needed to connect each node to the network.

Figure 5.2 shows a typical office layout for a small, five-member workgroup with room for expansion. The network

cabling is shown for reference only; the cabling in the example uses existing telephone wiring.

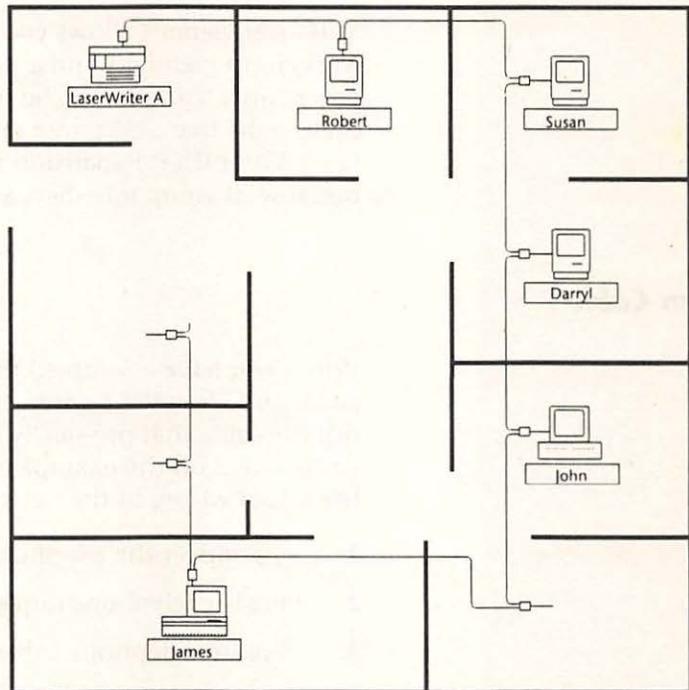


Figure 5.2 Typical office layout for a small workgroup.

Since the example uses existing telephone wiring for network services, each office will require only a PhoneNET connector and a standard 10-foot telephone extension cable. This will allow each workgroup member to position his or her computer virtually anywhere in their office.

Your total needs for the office shown in Figure 5.2 would be nine telephone extension cables, nine telephone duplex jacks, and nine PhoneNET connectors (eight DIN-8 connec-

tors for the Macintoshes and LaserWriter and one DB-9 connector for the IBM PS/2). You're also going to need a LocalTalk card for the single IBM PS/2 (shown as "John") in Figure 5.2.

This arrangement allows you to wire for a total of eight workgroup members and a printer. For the time being, you're only using six of the network connections: one for each of the five workgroup members and one for the shared LaserWriter IINT. Expansion room is available for two additional workgroup members and a file server in the future.

Run Cable

When you have assembled the cable, telephone duplex jacks, and PhoneNET connectors, you can begin to actually run the cable that physically connects each node to the network. Based on the example workgroup setting used so far, the actual wiring of the network is quite simple:

1. Disconnect the telephone cable from the wall jack.
2. Plug the telephone duplex connector into the wall jack.
3. Plug the telephone cable into the duplex connector.
4. Plug the network cable into the duplex connector.
5. Plug the other end of the network cable into the PhoneNET connector.
6. Plug the PhoneNET connector into the printer port of each Macintosh, and into the port on the IBM PS/2 LocalTalk card. The PhoneNET connector plugs into the LaserWriter IINT's AppleTalk port.

If your office telephone wiring precludes you from using the existing telephone wiring, now is the time to run the cable through the ceiling, cascading cable into each office. By cascading I mean that you have a cable coming *into* the office from the previous office and a cable going *out* to the next of-

fice. If you can't run the cable through the ceiling you'll have to call in a wiring specialist or—as a last resort—drill holes in the office walls. In any case, once you've got a network cable in each office, the cables plug into the PhoneNET connector and the PhoneNET connector plugs into the appropriate port as explained in step 5 and 6.

The first and last devices in a bus topology must be terminated. Farallon provides a terminator with each PhoneNET connector. The easiest way to deal with the termination issue is to locate the beginning and end of the network on your representative sketch of the network. In the case of the example used here, and shown in Figure 5.2, the "LaserWriter A" and the empty office node just below it would require termination. The terminator plugs into the PhoneNET connector in place of the second cable.

When you have each node physically connected to the network, it's wise to run a preliminary test on the network. You can use Technology Works' GraceLAN (discussed in Chapter 4: Groupware Implementations) to run this quick test following these steps:

1. Install the GraceLAN Responder startup document on each Macintosh on the network and restart them.
 - Make sure that each Mac on the network is running MultiFinder. This will be important in future steps of the network configuration process.
2. Launch the GraceLAN application.
3. Click on the Default Zone in the Zones window. A check mark will be displayed next to it, as shown in Figure 5.3

All devices should appear in a list in the appropriate window. Macintoshes will be listed in the Macintoshes window, Printers will be listed in the Printers window, and the PS/2 will be listed in the PCs window.

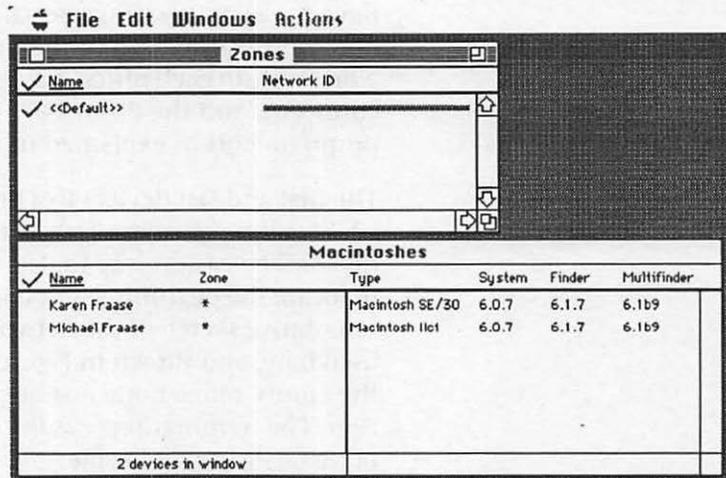


Figure 5.3 GraceLAN with the default zone selected.

If any device does not appear in the GraceLAN display, its connection is faulty. If this happens, check the node in question and restart the computer. Be sure to turn on AppleTalk in the Chooser desk accessory. It should then be displayed within the appropriate GraceLAN window.

While you've got GraceLAN up and running this wouldn't be a bad time to spot check the configuration status of all the Macintoshes on the network. A quick configuration display is provided for each Macintosh as shown in Figure 5.3. Check to be sure that each Macintosh is running the correct versions of all System Software components.

Establish File Sharing and Printer Service

When all the devices are properly connected to the network, you are ready to begin establishing the various network services that will be available to workgroup members. Begin

with establishing file service. Because your workgroup is relatively small, and you can safely expect average network traffic needs, a nondedicated file sharing system is the most appropriate choice.

Of the nondedicated file sharing alternatives, two stand out as excellent choices: Information Presentation Technologies Inc.'s Personal Server Network and the built-in file sharing capabilities of System 7.0. Both of these alternatives were discussed in Chapter 4: Groupware Implementations.

When you have installed and configured the file sharing software of your choice, run GraceLAN again to make sure that each of the workstations is configured properly. AppleShare Filing Protocol-compliant servers will be shown with the AppleShare server icon displayed next to them within the GraceLAN topology window, as shown in Figure 5.4.

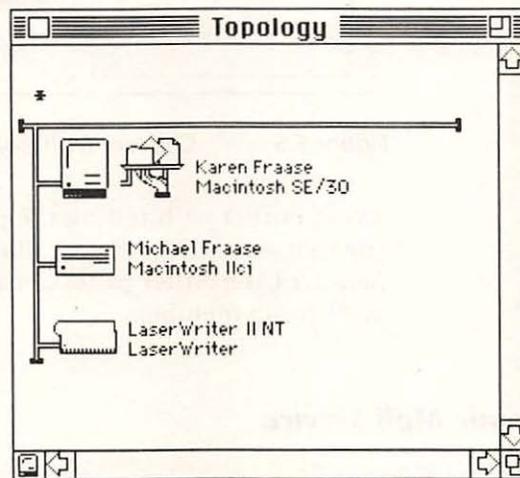


Figure 5.4 GraceLAN topology window.

The next step is to establish printer service for everyone connected to the network. Again, because this is a relatively

small workgroup with limited resources, the built-in print spooling available under MultiFinder or System 7.0 will be adequate. Activating print spooling is a simple matter of turning on Background Printing in the Chooser desk accessory as shown in Figure 5.5.

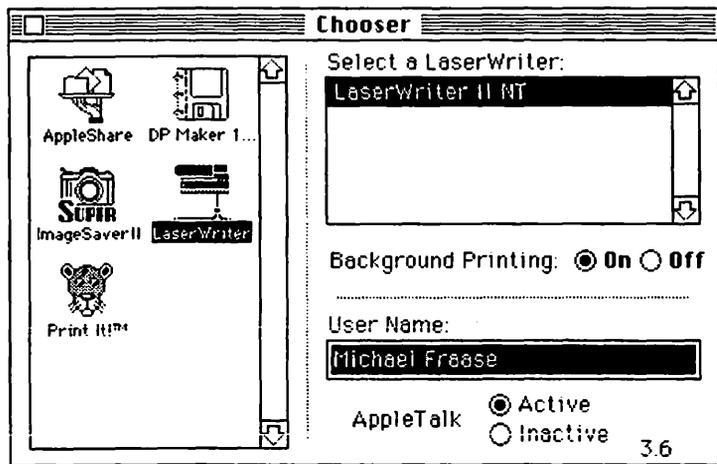


Figure 5.5 Chooser with Background Printing enabled.

Don't bother with testing the printer connection by printing to it across the network. If the printer is displayed in the Select a LaserWriter panel of the Chooser, it is available to workgroup members.

Establish Electronic Mail Service

The final basic network service to establish is electronic mail. For most workgroups CF Software's QuickMail is an excellent choice. QuickMail is the best alternative for a documentation workgroup, such as the one used for the examples here, because such workgroups frequently rely on freelancers for contract work. QuickMail has remote access

capabilities built-in. The QuickMail product is covered in Chapter 4: Groupware Implementations.

Begin the QuickMail installation procedure by deciding on the best Macintosh to use to provide electronic mail service to all workgroup members. The best choice would be the least used Macintosh with plenty of hard disk storage space. Follow the instructions in the QuickMail manuals to install the client software on each Macintosh. Install QuickMail/PC on the IBM PS/2. Finally, install the QuickMail MailCenter on the Macintosh designated as the electronic mail server.

Test Network

The final step in the network installation and configuration process is to quickly test everything. It is much too time-consuming to completely test everything; sending electronic mail from every workstation to every other workstation, for example, would take several hours.

A better approach is to run GraceLAN for a final time and take a close look at the topology window first to determine that it is complete in relation to your initial network diagram. Any faulty connections or configuration settings will result in the faulty node not appearing in the GraceLAN topology window.

Double-click on each network device to obtain detailed information about it. It's important to check and re-check all System Software versions, printer drivers, INITs, and CDEVs.

When you are satisfied that the network is properly installed and configured send electronic mail to each of the workgroup members explaining what you have done. Be sure to explain that the network is new and that problems will be unavoidable. Request each workgroup member's patience and explain the importance of reporting any problems that are encountered in using the network services. What you're

doing is basically asking the workgroup itself to provide the final testing for the network. Wait a few weeks until the network has stabilized before adding any other network services.

At this point it's crucial to explain the issues of general workspace security to the workgroup members. The best approach is to state—preferably in writing—that privacy will be respected and all attempts to preserve the privacy and integrity of individual workspaces will be made. It's important that everyone understands that accidents happen, and that privacy and security cannot be guaranteed regardless of good intentions.

It's important that everyone understand that rifling through someone's electronic workspace is just as offensive as rifling through someone's desk drawers. State clearly that such action will not be tolerated.

This might be a good time to provide each member with a data encryption program to allow private files to be archived in an encrypted state. This is also the time to explain the concept of access privileges on shared volumes and folders.

The idea of making volumes and files available to other workgroup members can be a little intimidating to some people. The access privileges feature of both approaches to file sharing covered here, however, allows each workgroup member to maintain control over which volumes, folders, and files are available to other workgroup members.

Each folder on each volume has an Owner, a Group, and a group of settings that allow the Owner to control which folders and files are available to other individuals or group members. Settings are available to disallow anyone from even *seeing* specified folders and their contents, let alone *accessing* them. Figure 5.6 shows the access privilege settings for the FrameMakerf folder residing on the Savannah server.

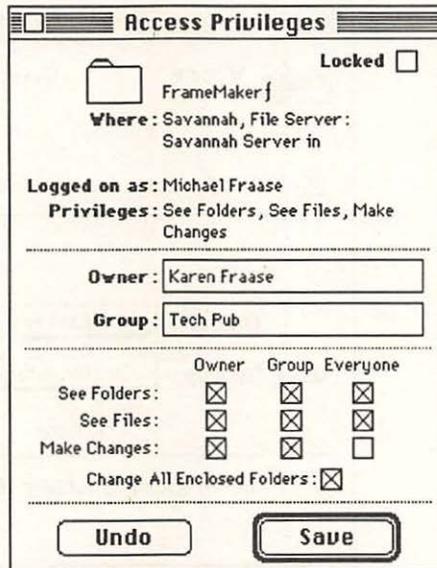


Figure 5.6 Access privilege settings for a remote folder.

As shown in Figure 5.6, the owner of the *FrameMaker f* on the Savannah server is Karen Fraase. She has set the access privileges for that folder (and all enclosed folders) so that everyone in the Tech Pub group can see the folders and files as well as make changes to the enclosed files. Any workgroup member outside of the Tech Pub group, however, would be prevented from making any changes to the *FrameMaker f* folder.

The file sharing feature of System 7.0 offers similar functionality for setting access privileges. Figure 5.7 shows access privilege settings for a local volume, making the contents of the volume named *Cassady* available to all workgroup members specified in the Documentation group.

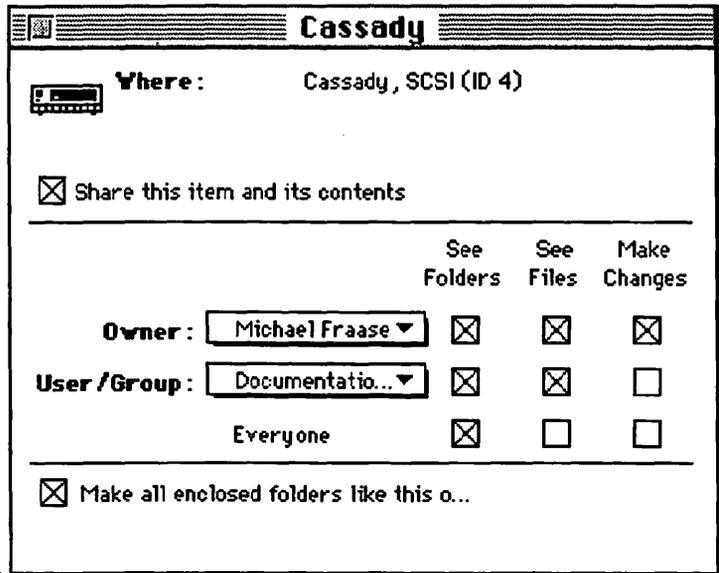


Figure 5.7 System 7.0 volume access privileges.

Using the Sharing command within the Finder lets each workgroup member define access privileges for entire volumes (as shown in Figure 5.7) as well as individual files. As shown in Figure 5.7 the owner of the volume has specified the following access privileges:

- The owner can see all folders, see all files, and make changes.
- Workgroup members of the Documentation group can see all folders and files, but cannot make changes.
- Anyone else within the workgroup can see all folders—but none of the files—and they are prevented from making changes.

Most network managers will find it best to advise each workgroup member to set the entire disk so that everyone can see

it, and then make specific folders private or assigned to defined groups.

Establish Screen-Sharing Service

After the network has stabilized, the next network service you may wish to add is screen-sharing. Farallon Computing's Timbuktu is a good choice for a workgroup of the type described here. Again, it's imperative that you *get permission from workgroup members to alter their personal workspaces*. A good approach is to send each workgroup member an electronic mail message explaining the benefits offered by the screen-sharing technology.

By now, some of the workgroup members may be sensitized to the privacy and security of their workspaces. It may be a wise alternative to place the Timbuktu files on an accessible file sharing volume or one of the file servers that is available to the entire workgroup. Instruct each member via electronic mail about how to install the software and let them do it themselves. Tell them that you're available to do the installation if they'd rather not be bothered with it. For those workgroup members that wish to do their own installation, make sure to provide them with the appropriate serial numbers via electronic mail.

It's important to explain the ramifications of the Timbuktu Guest Privileges settings to each workgroup member. Make sure that everyone understands that allowing certain Guest Privileges within Timbuktu overrides any access privileges defined on shared volumes. The exchange files setting, for example, refers to all files on the volume.

A Timbuktu Guest Privileges setting such as the one shown in Figure 5.8 would allow anyone with knowledge of the password "Group" to transfer *any* file on your volume, regardless of its file sharing access privileges setting.

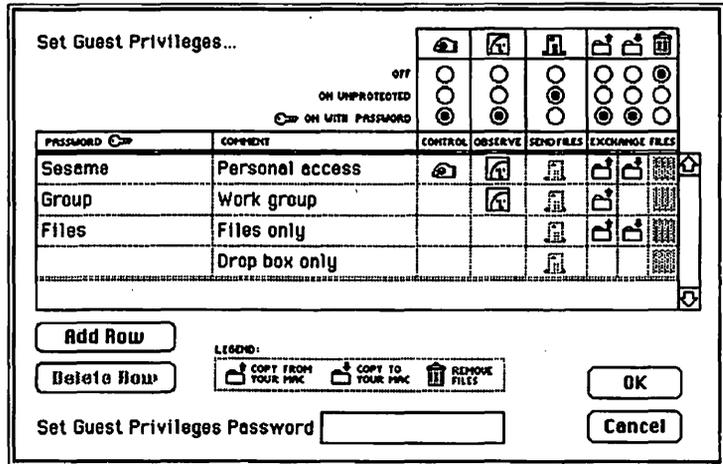


Figure 5.8 *Timbuktu Guest Privileges settings.*

Likewise, anyone with knowledge of the “Files” password would be able to retrieve any file from your hard disk as well as send any file to any folder within your workspace.

This situation is best approached by explaining the importance of maintaining proper security over one’s own workspace and by taking the time to explain the benefits of the screen-sharing technology.

Virtual Meeting Tutorial

Most collaborations are comprised of a series of meetings. Many of these meetings can be more effectively conducted in the virtual environment of the shared workspace. Most meetings are one of three types:

- *Presentations.* One workgroup member presents an idea or finding to the rest of the workgroup.

- *Brainstorming sessions.* Members of the workgroup meet to formulate a new idea or new approach to an existing problem.
- *One-on-one information exchanges.* Two workgroup members meet to share information with each other.

These meetings can all be handled effectively in the virtual shared space created when portions of all the private workspaces are linked via the local area network. This is not to say that the virtual shared space created by the workgroup is a replacement for all face-to-face meetings. There are times when traditional meetings are more appropriate and more effective. The virtual shared space will not take the place of traditional meetings, but it will enhance the workgroup's productivity by eliminating many "wasted" meetings.

A presentation meeting can be handled quite easily within the workgroup environment by using Timbuktu to display the presenter's screen in a window on all other workgroup members' monitors. The workgroup member giving the presentation would be the Timbuktu host, and all other workgroup members would log into his or her Macintosh as a Timbuktu guest.

The presenter would begin preparation for the meeting perhaps by sending electronic mail to each workgroup member needed to attend the presentation informing them of the time and date of the presentation. At this time the presenter would also forward to each appropriate workgroup member the password needed to log into the presenter's Timbuktu host at the time of the presentation. Each workgroup member would reply via electronic mail stating whether or not they would be available for the presentation. This is not absolutely necessary, but many workgroups may find that it helps ease the transition from physical to virtual meeting spaces.

Once the presentation time is determined, the workgroup member giving the presentation would perform the follow-

ing steps to register his or her Macintosh as a host and allow guest observation access:

1. Select Timbuktu from the Desk Accessory menu.
2. Select the Set Guest Privileges... command from the Timbuktu menu.
3. Specify the Guest Privileges for the presentation. If you will be doing presentations within the workgroup on a fairly regular basis, you may want to add a line to the Guest Privileges dialog box like the one labelled "Presentation" in Figure 5.9.
 - Click in the Observe column and a face icon will appear signifying that guests with the password assigned to this setting will be able to observe screen activity taking place on your Macintosh.
4. Assign a password allowing workgroup members to log onto your Timbuktu host as a guest in observe-only mode.
 - Notice that the highlighted setting in Figure 5.9 allows guests with the "Flipchart" password to log into your Timbuktu host in observe-only mode. Also notice that the drop box feature has been activated. This will allow workgroup members to send you comments on your presentation, but will not allow them to retrieve files from your drive or control your workspace.
5. Click the OK button.
6. Click the Guest Access On button in the main Timbuktu window. Your name will be automatically added to the list of hosts that are available to other workgroup members.

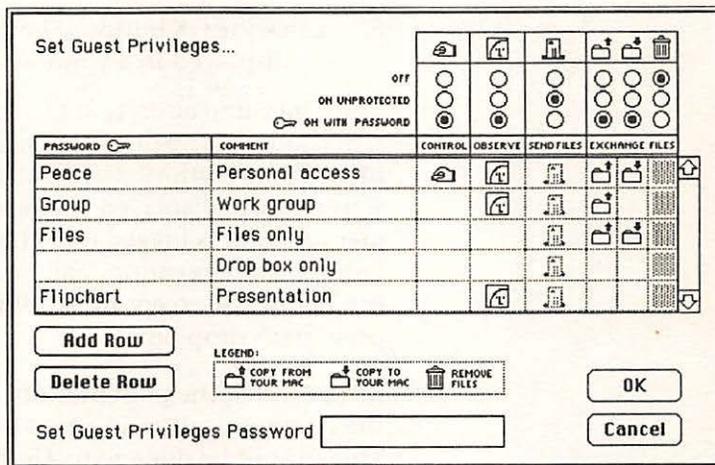


Figure 5.9 *Timbuktu Guest Privileges setting for an observe-only presentation.*

You may find that you want to do a practice run of your presentation with one other workgroup member as a Timbuktu guest. This can be done without defining a special password or changing any settings by selecting the Wait For Guest command from the Timbuktu menu.

Workgroup members that are attending the presentation do not have to alter their Timbuktu settings in any way. All that's required is to remember the password that was assigned by the presenter. At the time of the scheduled presentation, each workgroup member that will be participating in the meeting follows these steps to log into the presenter's Timbuktu host:

1. Select Timbuktu from the Desk Accessory menu.
2. Select the appropriate host in the host list by clicking on the name to highlight it.
3. Click the Observe button.

4. Enter the password that was assigned by the presenter.
5. Click the OK button. The presenter's screen activity will be displayed in a window on your monitor.

The communication that takes place in a screen-sharing presentation like the one described here is a one-way communication: everything that takes place on the presenter's screen is also displayed on the monitor of every workgroup member that is logged into the Timbuktu host. The workgroup members cannot communicate with the presenter using Timbuktu except by sending comments or files to the presenter's drop box.

At the end of the presentation, the host may decide to open the communication channel to two-way communications. This cannot be done with Timbuktu, at least not with the settings defined for the presentation. The workgroup does have another collaborative computing tool close at hand, however.

The conferencing feature of QuickMail can be used for an informal roundtable discussion at the close of a presentation like the one outlined in this example. An example of using this technique is provided in the next section.

Timbuktu is also useful for simulating the other two common meeting types, brainstorming sessions and one-on-one information exchanges. Examples of both of these will be provided in later sections.

Computer Conferencing Tutorial

QuickMail's conferencing functionality is a feature that is easily underrated. In the previous section, an example of using Timbuktu to simulate a presentation-type meeting was provided. Presentation meetings usually end with some sort of roundtable discussion or broad information exchange between participants. QuickMail's conferencing feature can be

used in conjunction with Timbuktu to provide this round-table discussion capability. The QuickMail conferencing feature can also be used by itself to facilitate impromptu conferences and brief virtual meetings.

To begin a conference, any one of the workgroup members can initiate the dialogue with the following steps:

1. Select QuickMail from the Desk Accessory menu.
2. Select the Conference command from the QuickMail menu, or use the Command-T keyboard equivalent. QuickMail's conference window will be displayed.
 - All available conferees will be listed in the Online Users panel.
 - What you type will be displayed in the Entry panel.
 - A complete transcript of the conference will be displayed in the Transcript panel.
 - Across the top of the conference window are buttons that enable you to switch zones, send the contents of the Entry panel, save the conference transcript in a text file, and save the conference transcript on the Clipboard
3. Select the conferees you want to broadcast a text message to by selecting them in the Online Users panel. Shift-clicking within the list allows any number of conferees to be selected.
4. Type your message in the Entry panel.
5. Click the Send button or press Enter to broadcast your text message to the selected conferees.
 - The QuickMail conference window will automatically open on each conferees Macintosh. They can respond to your message by the steps outlined above.

- You can send private, or semi-private messages to any conferee by selecting only the desired recipients in the Online Users panel and sending your text message as usual.
- You can lock out the conferencing feature of QuickMail by using the Privacy command on the QuickMail menu.

Figure 5.10 shows the QuickMail conference window with an active two-person conference.

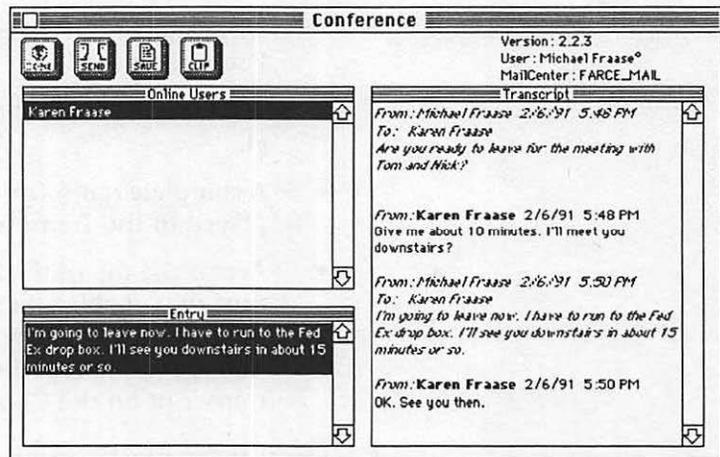


Figure 5.10 QuickMail conference window.

It's important to use some common sense when using the conference facility of QuickMail or conferences will quickly devolve into meaningless chaos. Even though QuickMail allows you to confer with as many workgroup members as you like—even across network zones—smaller conferences work best. Even in small conferences, you should avoid questions that generate useless responses. Asking an entire group if they all understand what's been presented is too open-ended and will result in a flood of meaningless replies. If some-

one doesn't understand something that's been covered, trust them to let you know in a more appropriate manner.

The conferencing feature of QuickMail can also be used to communicate quickly with individuals on the network. It's easier and faster, for example, to use the Conference command than dialing an inter-office extension on the telephone.

Teleconferencing Tutorial

CE Software's QuickMail provides excellent electronic mail services for the workgroup, and it even offers a public bulletin board that can be used to post messages to all workgroup members. QuickMail was designed as an electronic mail system, though, not a teleconferencing system.

Many collaborative workgroups can benefit from ongoing discussion groups that are relevant to their areas of expertise. The documentation group used in this chapter as an example, for instance, would be interested in following and discussing developments ranging from online hypertext systems to groupware and perhaps even virtual reality.

An electronic mail system is not useful for ongoing discussion groups like this, because the focus is on one-to-one or one-to-many communication, not the many-to-many communication that typifies the traditional discussion group.

Spider Island Software's excellent teleconferencing application, TeleFinder, was discussed in Chapter Four: Groupware Implementations. TeleFinder provides easy to use teleconferencing for remote workgroup members, but in its current version does nothing for local area network teleconferencing. The good news is that Spider Island Software is in the process of modifying TeleFinder to work as both a remote and local teleconferencing system. Until then, however, there is a workaround following these steps:

1. Install and configure TeleFinder as usual, placing the TeleFinder conferences and library file folders on a volume that is accessible to workgroup members.
2. Set the TeleFinder conferences folder access privileges to allow all workgroup members to see and change the contents of the folder.
 - Remote workgroup members will access TeleFinder via modem as usual.
 - Local workgroup members can access the TeleFinder conferences and library file folders across the network as a shared folder.

Eventually, TeleFinder will work in both local and remote modes. It will accomplish this by using Apple's Communications Toolbox and a custom startup document to allow TeleFinder to be used across a local area network. In all other respects it will operate identically as it does in its remote mode.

Project Management Tutorial

In order for any collaborative workgroup to be viable, it must manage projects and billing effectively. The documentation workgroup used as an example in this chapter, for instance, must track manual projects, online help system development, illustration tasks, and client billing. Even if the workgroup is part of a larger enterprise, an effective billing system must be implemented because most corporate business units actually bill for their services within the enterprise. This will become even more common as large companies make the shift from a vertical hierarchy to flatter, horizontal organizations with independent business units.

Projects can be managed effectively using any of the commercially available software packages such as MacProject II from Claris Corporation. None of the currently available

project management applications are designed for multiple users. This usually won't pose much of a problem because there's usually a single person ultimately responsible for meeting milestones and deadlines of any project. In the case of the documentation workgroup used as an example in this chapter, an editor or lead writer would be responsible for project management.

Billing management is another issue completely. In the documentation workgroup, for example, each writer, editor, illustrator, and designer must keep track of his or her billable hours on a project-by-project basis. Most small workgroups will be involved in more than a single project at a time. In most situations the collecting and compilation of workgroup members' timeslips and generation of invoices would fall on the office manager or equivalent. It's a time-consuming job, especially in workgroups. Some workgroups may find it advantageous to hire an associate just to manage this time-consuming task.

Most workgroups, however, can oversee their own billing management tasks with the Network Edition of Timeslips III from Essex, Massachusetts-based Timeslips Corporation. Timeslips uses a familiar slip entry form for time and expense tracking and charges can be generated by an hourly rate or a flat fee. Invoices can then be generated from a compilation of all workgroup members' timeslips. A variety of management reports including an accounts receivable aging report can also be generated.

Timeslips installs as a desk accessory and associated files on each workgroup member's computer (a version is also available for the IBM PC). The program works in much the same way as paper-based time-tracking systems that generate a slip for each task and client, with each workgroup member maintaining a series of electronic timeslips for their services.

The billing management system is set up by the network manager who defines a set of clients, hourly rates for each

workgroup member, and a list of projects currently under proposal or in process.

When Timeslips has been installed on each computer and properly configured for the workgroup, each workgroup member uses the Timeslips desk accessory to maintain an electronic record of services performed, time spent on various projects, and expenses.

To start a new timeslip you select TSTimer from the Desk Accessory menu or launch the TSTimer application. This would cause Timeslips to display the last timeslip that was active. If this is the start of a new billing period or a new installation, the program would alert you that no timeslips exist and that a new one is being created. Selecting the New Timeslip command from the Timeslips menu would create a new Timeslip similar to the one shown in Figure 5.11.

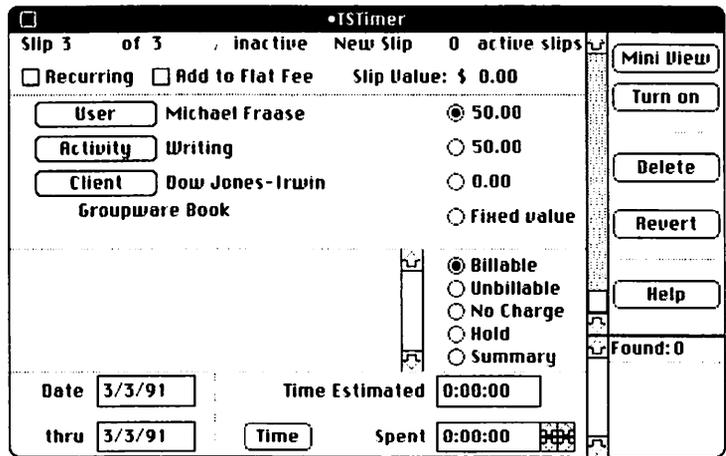


Figure 5.11 Timeslips window displaying a new timeslip.

The New Timeslip command creates a new timeslip with a value of \$0.00 and it copies the first four fields (User, Activity, Client, and Matter) from the previous slip if one is avail-

able. If not, all fields are blank except for the two date fields at the bottom left panel of the timeslip. To modify the new timeslip for use you follow these steps:

1. Click the User button to assign the correct user to the new timeslip.
2. Select the correct user from the dialog box that is displayed.
3. Click the Activity button to assign the correct activity to the new timeslip.
4. Select the appropriate activity from the dialog box that is displayed.
5. Click the Client button to assign the correct client and matter to the new timeslip.
6. Select the appropriate client to be billed from the dialog box that is displayed.
 - Click the Matter button in the dialog box and select the appropriate matter for the timeslip.
7. Enter a note indicating the service performed and any other notation necessary for the timeslip in the scrolling text field provided.
8. Enter the time you estimate the task to take in the Time Estimated field.
9. Click the Turn on button to start the stopwatch built in to Timeslips.
 - Alternatively, you can enter the actual number of hours spent on the task in the Spent field instead of turning on Timeslips internal stopwatch.
10. When you are finished with the task, click the Turn off button to stop the clock and log the time spent on the task.

- Alternatively, you can enter the actual number of hours spent on the task in the Spent field, bypassing the internal elapsed time clock.

The completed timeslip will look similar to the one shown in Figure 5.12.

The screenshot shows a window titled "TSTimer" with the following details:

- Slip 3 of 3 inactive, 0 active slips
- Recurring: Add to Flat Fee: Slip Value: \$ 400.00
- User: Michael Fraase (radio button selected at 50.00)
- Activity: Writing (radio button at 50.00)
- Client: Dow Jones-Irwin (radio button at 0.00)
- Groupware Book (radio button at Fixed value)
- Tutorials: 10:00 a.m. - 6:00 p.m.
- Billable: Unbillable: No Charge: Hold: Summary:
- Date: 3/3/91 Time Estimated: 10:00:00
- thru: 3/3/91 Time Spent: 8:00:00
- Buttons: Mini View, Turn on, Delete, Revert, Help, Found: 0

Figure 5.12 Timeslips window displaying a completed timeslip.

Each workgroup member is responsible for maintaining his or her timeslips and expense slips throughout the billing period. The Timeslips data files are shared by all workgroup members and the files themselves reside on the file server.

At the end of each billing cycle the person responsible for generating the invoices uses the second half of the Timeslips software, the TSReport application, to generate bills and perform various analysis on the billing and expense information. The following tutorial assumes that all client and invoice format information is configured properly and up to date.

To generate invoices, follow these steps:

1. Launch the TSReport application.
 - You'll be prompted to make sure that the date information is correct. This is critical for accounts receivable aging and interest calculations.
2. Select the Transactions... command from the Clients menu to update any payments that may not have been entered into the billing system. A window like the one in Figure 5.13 will be displayed for each client.

Date	Description	Amount	Qty	Adj%	Tax
3/3/91	Payment	0.00			

Figure 5.13 Timeslips Transactions window.

- Enter the amount received in the appropriate field and click the Accept button.
3. Select the Bills... command from the Reports menu. A window like the one shown in Figure 5.14 will be displayed.

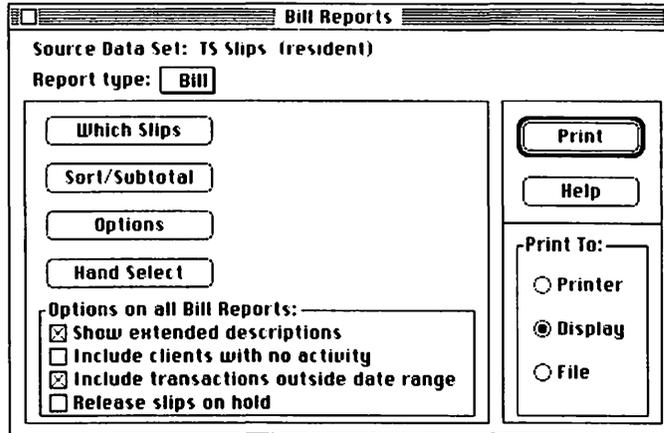


Figure 5.14 Timeslips Bill Reports window.

4. Click the Which Slips button to select the timeslips to be included in this billing cycle. The default is all timeslips that have not already been billed. The selection criteria dialog box is shown in Figure 5.15.

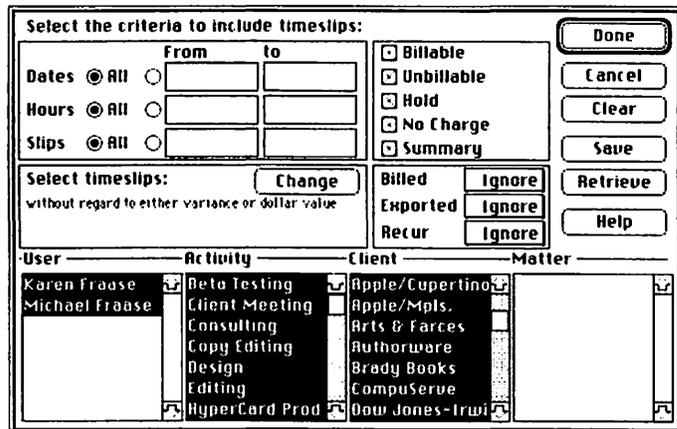


Figure 5.15 Timeslips bill selection criteria dialog box.

- If you want to print only a single bill, override the default to select the appropriate client. This will restrict the retrieved timeslips to only that client.
5. Click the Done button to return to the Bill Reports window.
 6. Click the Display radio button to direct a printed preview of the bills to the screen.
 7. Check to make sure that the bill appears properly.
 8. Click the Printer radio button to direct output to the printer.
 9. Click the Print button to actually print the bills.
 - When the bills have been printed, a dialog box will appear asking if you want to finalize the billing. If you are absolutely sure that the bills are correct, click the Backup and Finalize button. This will backup the current data set to a different file and finalize the billing. If you discover that the billing is incorrect, you can restore the data set from the backup.

This sequence is repeated at the end of each billing cycle. In addition to printed invoices, Timeslips also provides the workgroup with a variety of printed reports. Reports are available for billing details, summaries, graphs of billable hours per workgroup member and income, and the like.

A Collaborative Session Using Inspiration with Timbuktu

One of the most effective methods of collaboration is the brainstorming session. A brainstorming session can be made even more effective if two or more workgroup members can work interactively on the same document. This is possible to do with the collaborative environment used as an example in this session. Timbuktu can be used to add interactivity be-

tween two or more workgroup members to Inspiration, a unique thought processor published by Ceres Software of Portland, Oregon.

Inspiration is a tool for working with ideas and information that enables you to work simultaneously with both visual diagrams and text by providing two views of the same information: outlines and diagrams. Inspiration can be thought of as a sort of whole-brain outliner. It uses the visual, right-brain to form ideas and think visually and the logical, left-brain to organize thoughts into outlines.

For the documentation workgroup used as the example in this chapter, Inspiration is a very useful addition and can be used as a brainstorming tool by virtually any workgroup. When used in conjunction with Farallon's Timbuktu, the power of Inspiration is enhanced significantly by allowing ideas to be outlined and diagrammed interactively, across the network.

This tutorial will provide step-by-step instructions for using the Inspiration/Timbuktu combination to collaborate on the underlying concepts for a publication proposal under consideration by the documentation workgroup. It's assumed that Timbuktu and Inspiration are installed and properly configured on each workgroup member's Macintosh.

Begin by setting a time for the collaborative Inspiration session using QuickMail. The electronic mail message should detail the time of the session, designate who will act as the Timbuktu host for the session, provide an overview of the topic, and specify a Timbuktu guest password for the session. Just before the appointed time for the interactive session, the host should perform these steps to prepare his or her workspace for the collaborative guests:

1. Select Timbuktu from the Desk Accessory menu.
2. Select the Set Guest Privileges... command from the Timbuktu menu.
3. Specify the Guest Privileges for the presentation. If you will be doing presentations within the workgroup on a fairly regular basis, you may want to add a line to the Guest Privileges dialog box like the one labelled "Inspiration" in Figure 5.16.
 - Click in the Control column and a face icon will appear signifying that guests with the password assigned to this setting will be able to control your Macintosh.
 - It's important to note that when using Timbuktu in this manner, the host maintains ultimate control over his or her workspace, and any action performed by the host takes precedence over actions performed by guests.
4. Assign a password allowing workgroup members to log onto your Timbuktu host as a guest in control mode.
 - Notice that the highlighted setting in Figure 5.16 allows guests with the "WholeBrain" password to log into your Timbuktu host in control mode. Also notice that the drop box feature has been activated. This will allow workgroup members to send you comments on the interactive session while it is active, but will not allow them to retrieve files from your drive.
5. Click the OK button.
6. Click the Guest Access On button in the main Timbuktu window. Your name will be automatically added to the list of hosts that are available to other workgroup members.

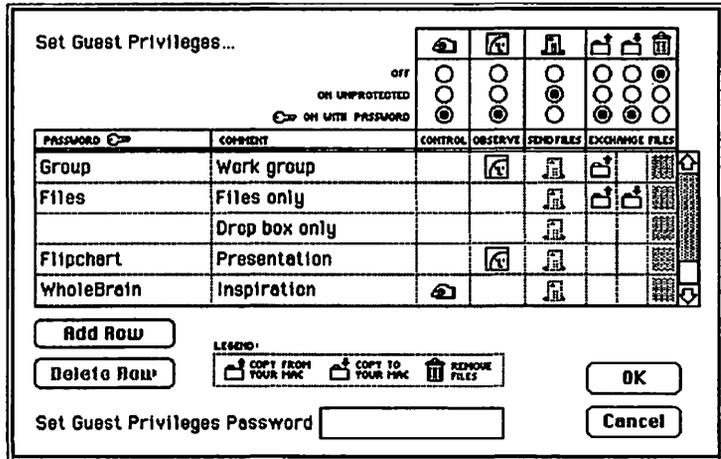


Figure 5.16 *Timbuktu Guest Privileges setting for an interactive session.*

7. Launch the Inspiration application.
8. Wait for the workgroup members to log into your Timbuktu host.

Workgroup members that are participating in the interactive session do not have to alter their Timbuktu settings in any way. All that's required is to remember the password that was assigned by the host. At the time of the scheduled session, each workgroup member that will be participating in the meeting follows these steps to log into the designated Timbuktu host:

1. Select Timbuktu from the Desk Accessory menu.
2. Select the appropriate host in the host list by clicking on the name to highlight it.
3. Click the Control button.
4. Enter the password that was assigned by the host.

5. Click the OK button. The host's screen activity will be displayed in a window on your monitor. Any action that you take within the Timbuktu window will also be displayed within the shared space.

The communication that takes place in a screen-sharing collaboration like the one described here is a two-way communication: everything that takes place on the host's screen is also displayed on the monitor of every workgroup member that is logged into the Timbuktu host. Similarly, any action taken by any workgroup member is automatically updated within the shared space.

The interactive session will open with an untitled Inspiration document being displayed on everyone's screen. The Timbuktu host with the Inspiration document will appear in a window on each guest's Macintosh. An example is shown in Figure 5.17.

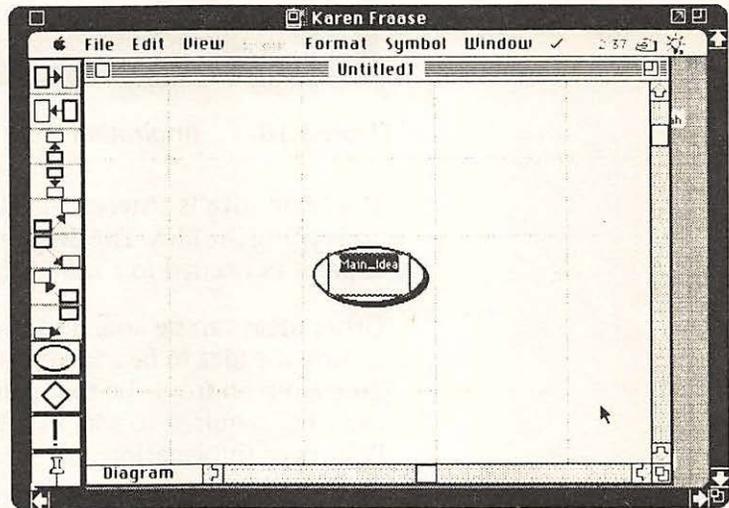


Figure 5.17 Inspiration document accessible in Timbuktu window.

When all the workgroup members are logged into the Timbuktu host, the host should begin the interactive session by beginning to add to the Inspiration diagram, by adding the Main Idea. An example is shown in Figure 5.18.

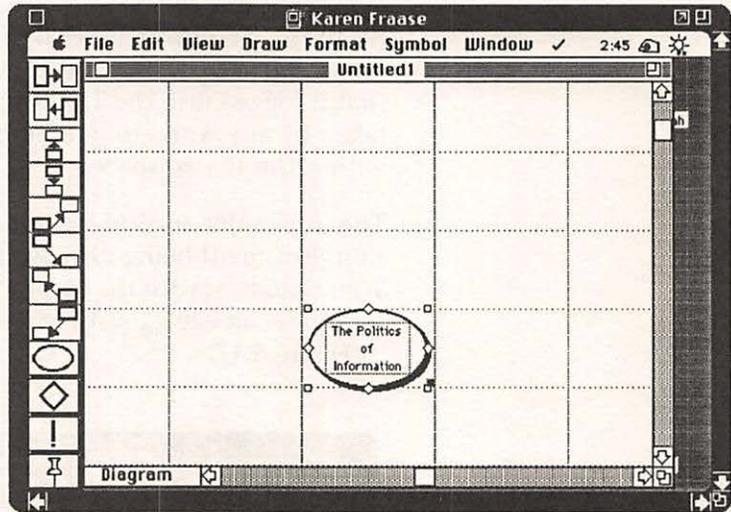


Figure 5.18 *Inspiration Main Idea in Timbuktu window.*

The Main Idea is entered by selecting the diagram element and typing the idea. The diagram element will automatically expand as needed to contain the information.

Other ideas can be added to the Inspiration diagram by selecting the idea to be added to and selecting the appropriate diagram icon from the tool palette. For example, here are the steps required to add a subtopic to the Main Idea: The Politics of Information:

1. Select The Politics of Information diagram element as shown in Figure 5.18.
2. Click the Left Subtopic diagram icon (eighth icon from the top) on the tool palette.

- A new subtopic diagram element will appear below and to the left of the Main Idea.
3. Enter the subtopic idea. The diagram element will expand as needed to contain the information.
- Figure 5.18 shows the result of adding a subtopic within the shared workspace.

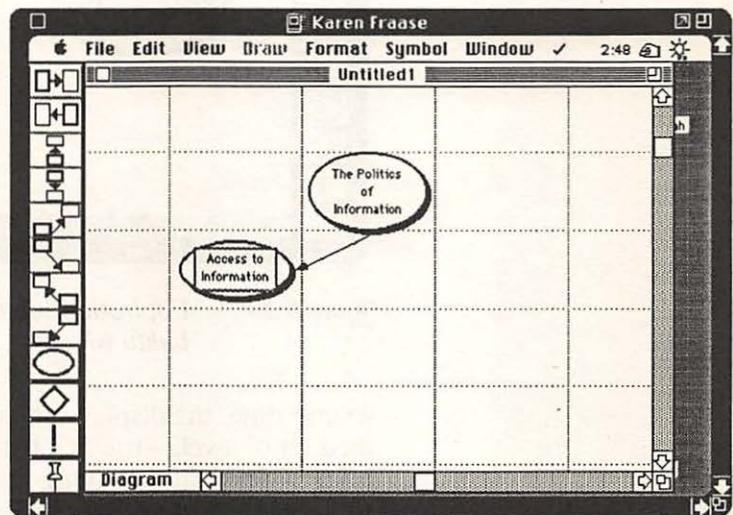


Figure 5.19 *Inspiration Main Idea with subtopic in Timbuktu window.*

Any number of subtopics and sub-subtopics can be added in the same way until the group runs out of ideas. Diagram elements can be moved around simply by selecting them and dragging to the new position.

In addition, labels can be applied to the various diagram elements to help the workgroup visualize topics and subtopics by using any of the commands available on the Symbol menu. Figure 5.20 shows an additional subtopic and a set of visual labels.

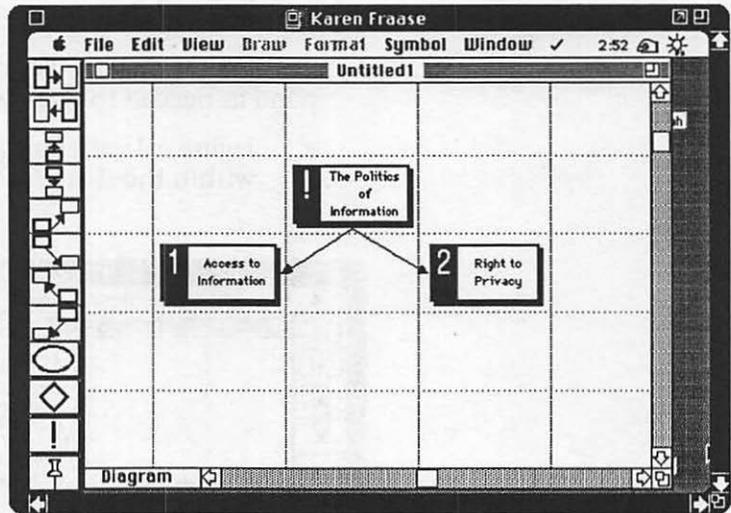


Figure 5.20 *Inspiration subtopics and visual labels in Timbuktu window.*

At any time, the display can be revised to show only a specified set of levels—the first three levels of subtopics, for instance—without damaging or altering the underlying subtopics of their structure within the workspace. These views can be toggled by any member of the workgroup, although it is usually best if all views on the workspace are controlled by the host or facilitator to prevent confusion and frustration.

Some members of the workgroup may think more logically and would be more comfortable with an outline view of the information contained in the shared workspace. At any time, the display can be changed to reflect an outline view of the material by using the Outline command on the View menu. This assures that all workgroup members are able to contribute to the collaborative session, and everyone feels heard. Figure 5.21 shows an outline view of the shared workspace in the same state as that of Figure 5.20.

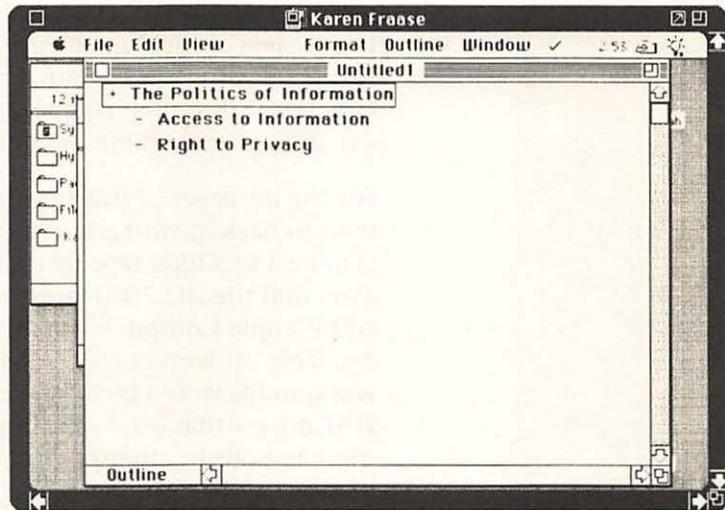


Figure 5.21 Inspiration outline view in Timbuktu window.

When the workgroup runs out of ideas, the topics and sub-topics can be organized by priority by selecting the item and dragging it to a new location. This prioritizing can take place in either the diagram or outline views, so again, everyone is able to contribute to the process and feel heard.

The final Inspiration document can be placed on the file server for access by all members of the workgroup at their leisure. Each member can refer to the shared document, or make copies of it, as he or she sees fit.

Workgroup Data Backup Tutorial

Perhaps the most mundane task of collaborative computing is assuring that the individual workspaces of each workgroup member are properly backed up. Berkeley, California-based Dantz Development Corporation offer the most ele-

gant solution to this task that is often overlooked in a workgroup environment. Dantz's Retrospect Remote automates the process of backing up any storage device on the network to a wide variety of backup devices including digital audio tape (DAT), 8mm videotape, standard tape cartridges, optical discs, or removable cartridges.

For the purposes of this tutorial, Retrospect Remote will be used to backup workgroup members hard disk drives to a standard DC-2000 tape drive. It's important to realize, however, that the DC-2000 tape drives—such as the one marketed by Apple Computer—are the least practical option available for workgroups. A better alternative for almost all workgroups would be larger capacity tape drives—such as a DAT drive—that offer significantly faster throughput. It's quite tedious to attempt to regularly backup the enormous amount of data produced by a workgroup to a series of tape cartridges that can hold less than 40 MBytes each.

Retrospect Remote is capable of backing up any volume that can be mounted on the local desktop. In addition, it can be used to backup any remote volume with the addition of a startup document to each workgroup member's System Folder.

Explain to the members of the workgroup the importance of maintaining a backup everyone's workspace. Sooner or later, an error will occur that will cause a loss of data. With a backup, lost work is limited to only what was done since the last time the member's workspace was backed up.

In order to maintain backups of everyone's workspace, however, Retrospect Remote (and by extension the person responsible for performing the actual backup activities) will have complete access to *all* data contained within the workspace. Be sure to inform each workgroup member that this is necessary to maintain the integrity of the workgroup and that personal or sensitive data can be encrypted. Retrospect also offers the option of restricting access to any folder by renaming any sensitive folder with a bullet (•) preceding the

folder name. In such situations folders marked as restricted will not be backed up. It is *absolutely essential* to get permission from workgroup members before altering or accessing their personal workspaces.

When you have explained the backup procedure to all the workgroup members and have obtained their permission to access their workspaces install the Remote startup document on each workstation as described in the Retrospect Remote manual. Explain to the workgroup members how to restrict backup access to their workspaces and show them how to configure the Remote startup document via the Control Panel.

Next, designate a Macintosh to perform the backup tasks for the entire workgroup. This can be any Macintosh on the network with one or more backup devices connected that is running the Retrospect Remote application.

The final step of the installation and configuration procedure is to activate all the Remote startup documents by following these steps:

1. Launch the Retrospect application.
2. Select the Remotes command from the Config menu. A list of Known Remotes will be displayed.
3. Click the Network... button. The Known Remotes list will be updated to include all Macintoshes that are running the Remote startup document.
4. Select any Macintosh in the list and click the Install button.
5. Enter an Activator code (packaged with Retrospect) in the dialog and click the OK button.
6. Continue these steps for each workstation.

Retrospect offers advanced scheduling and scripting capabilities, but if your workgroup is like most, your main concern isn't in scheduling backup sessions just yet. You're more

concerned about getting everybody backed up as soon and as quickly as possible. Here's the quickest way to do it:

1. Launch the Retrospect application. Retrospect's main window will be displayed, as shown in Figure 5.22.

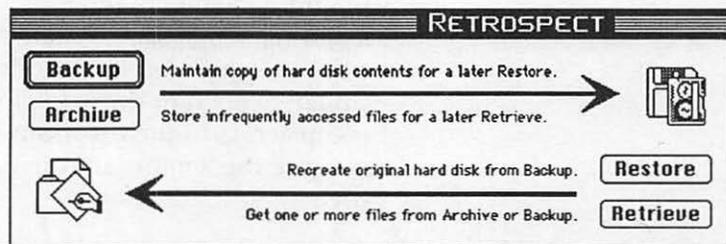


Figure 5.22 Retrospect main activity window.

2. Click the Backup button. Retrospect's Backup window will be displayed with a list of available sources for backup, as shown in Figure 5.23.

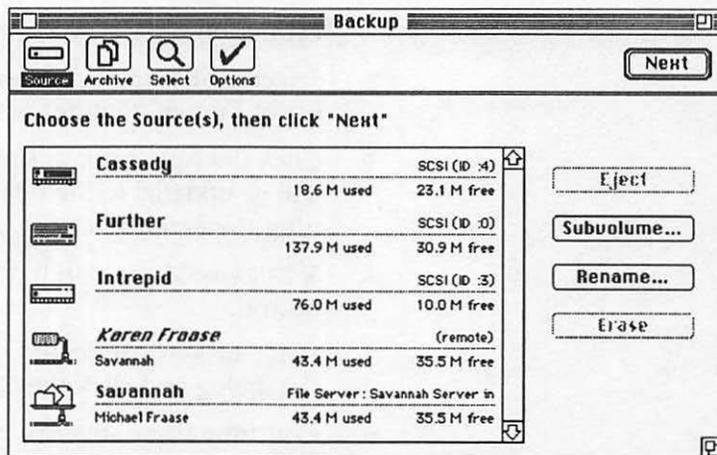


Figure 5.23 Retrospect Backup window.

3. Select the source volume you want to backup and click the Next button. Retrospect's Archive window will be displayed with a list of available archives. Since this is an initial backup, no archives will be available.
4. Click the Create New... button. Retrospect will prompt you to choose an archive type.
5. Click the DC-2000 option in the devices list and click the Next button. You'll be prompted to name and save the new archive.
6. Select the newly created archive in the Archive window and click the Next button. Retrospect will scan the source volume and compare it with the selected archive. Since it's a new archive, all files will be selected and displayed in Retrospect's selection window, as shown in Figure 5.24

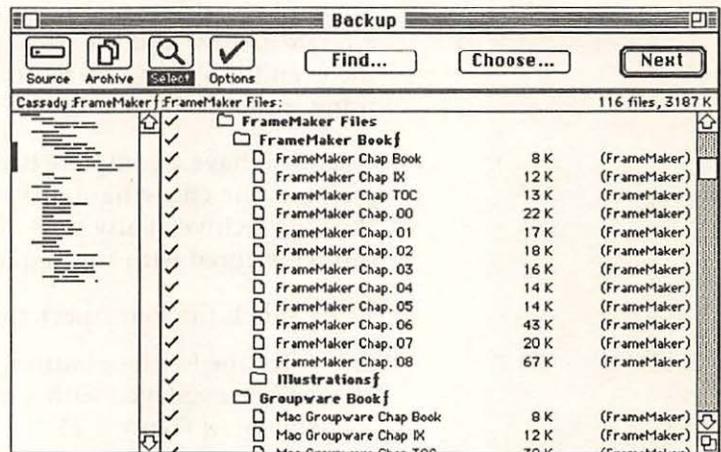


Figure 5.24 Retrospect file selection window.

7. Click the Next button. Retrospect's Option window will be displayed allowing you to turn on verification and compression.

8. Check the Compress files checkbox.
9. Click the Execute Now button. Retrospect will prompt you to insert the tape cartridge that contains the selected archive. Since this is a new archive, insert a new tape.
10. Retrospect will prompt you to select the destination media for the backup. Select the tape cartridge you have inserted.
 - If more than one tape cartridge is needed, Retrospect will prompt you at the correct times to insert a new tape cartridge.

As mentioned earlier, the larger capacity the backup device the better for most workgroup situations. Given the DC-2000 tape drive used in this tutorial, you would have to repeat the above steps for each hard disk drive on the network.

If you used a more advanced device, such as a DAT drive, all the hard disk drives on the network could be combined into a single archive and the backup started before you left for the evening. When you returned the next morning, everything would be backed up.

Once you have a complete backup for each workgroup member, the entire hard disk drive can be recovered from the tape archive at any time. Additionally, individual files can be restored with this sequence:

1. Launch the Retrospect application.
2. Click the Retrieve button. Retrospect's Archive window will be displayed with a list of available archives, as shown in Figure 5.25.

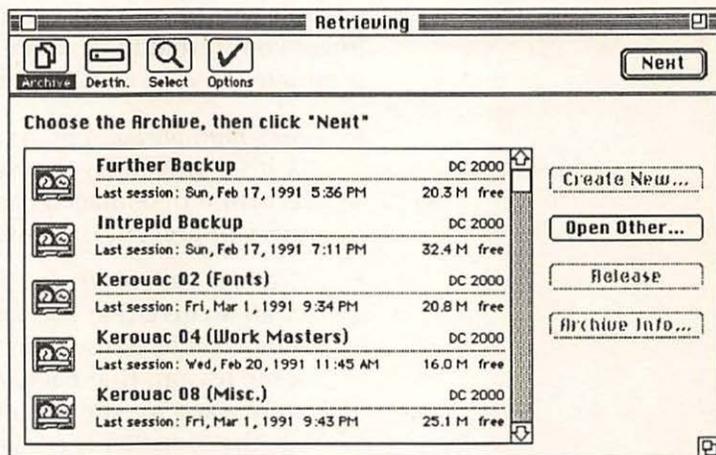


Figure 5.25 Retrospect Archive window.

3. Select the Archive you want to retrieve files from and click the Next button.
4. Select the Destination volume you want to recover the files to and click the Next button.
5. Select the files you want to retrieve from the list and click the Next button. Retrospect's file retrieval Options will be displayed.
6. Select the options you want in the Options window and click the Execute Now button. The files you specified will be retrieved using the options you selected.

It's important to keep a separate backup of each archive catalog for all work disks within the workgroup. Retrospect can re-catalog any archive, but the process—at least with DC-2000 tape cartridges—is painfully slow. It takes almost as long to re-catalog the archive as it does to actually restore the information. This is not a flaw in Retrospect, but rather a nature of the tape backup medium.

It's important that you develop a backup strategy as soon as possible within your workgroup. Here are three general rules to get you started in developing a strategy that is most appropriate for your working environment:

- *Keep multiple backups.* As reliable as most backup media is, it's highly likely that sooner or later you'll encounter a corrupt or damaged backup or archive. Multiple backups will prevent the loss of any data.
- *Keep one backup off-site.* If you experience a media error or an archive that has somehow been corrupted, your reflex will be to use one of your multiple backups. If for some reason, that back up is unusable, you'll immediately reach for the remaining archive and if you damage that one you're really out of luck. Also, in the event of a burglary, fire, or other calamity your workgroup's data will be retrievable if a backup is stored off-site.
- *Recycle archive media.* There's little reason to keep an infinite number of backups or archives. Especially expensive media like optical discs or removable cartridges can be effectively recycled by developing an appropriate backup strategy.

You'll find that Retrospect will help you develop an effective backup strategy with its extensive scripting and scheduling capabilities.

Chapter Six

Creating the Collaborative Environment

People hate meetings. Most people perceive meetings as a mixture of boredom and frustration, where nothing much gets done. People want to be productive, but the typical meeting has come to be perceived as counterproductive. Most meetings are indeed counterproductive because they tend to take place in counterproductive environments.

The typical meeting takes place with six to ten people sitting around a table. The meeting leader is at the head of the table with visual aids ranging from a whiteboard and flip charts to an overhead projector or full-blown multimedia presentation. The individuals take turns talking and frustration mounts in direct proportion to the amount of monopolization and interruption that takes place. Each time a new speaker takes the floor, all the other members' attention shifts to the new speaker. Most of the group looks to the leader for reaction cues.

For whatever reasons, the meeting has come to signify collaboration in most of our organizations. Michael Schrage has identified one of the main reasons why many meetings don't work well. Traditional meetings are structured in a way that enables individuals to make their points, but prevents the group from creating a shared understanding. "The typi-

cal meeting environment subsidizes the excesses of individual communication at the expense of collaborative community," according to Schrage.

Schrage proposes changing a meeting's ecology as a solution to this set of problems. By changing a meeting's ecology, individual behavior shifts from ego-centric to focusing on the group. "The environment is far more responsive than that of a traditionally structured meeting. Many of the behaviors associated with typical meetings either disappear or are transformed," notes Schrage. "Long-winded soliloquies become short gusts. The quantity of speech is different; the quality of speech is different. The individual ego evolves into a collective ego. The focus is no longer on who said what but on what's on the screen. It's not that comments are anonymous, it's that presenting them on the screen physically decouples them from the speaker. The comment literally becomes a contribution to the meeting. Others are free to add to or build on the comment as they see fit."

The best way to enable effective group collaboration is to provide an appropriate environment, appropriate tools, and changing the ecology of the meeting structure. With that in mind, this chapter discusses how to go about providing those things and making the necessary changes.

Removing Inhibitions to Collaboration

In order to create effective environments for collaboration, we must change the ecology of a meeting. Michael Schrage has defined the ecology of a meeting as the "interplay of ideas, personalities, and environment." Changing the ecology of a meeting allows the collaborators, and especially the person responsible for the meeting's outcome, to forget about trying to change individual behavior.

The key to changing the ecology of a meeting is to remove the presentation aids that are static and replace them with a

shared space that is dynamic and accessible to all members of the meeting.

According to Michael Doyle and David Straus in *How to Make Meetings Work*, all meeting ecologies are comprised of two distinct aspects: content and process. Content is the actual messages that are exchanged within the meeting. Process is the way in which the messages are communicated. Doyle and Straus recognized a tendency in the person running the meeting, if she or he is process-oriented (most are), to attend to the process issues at the expense of active participation by other members.

Doyle and Straus advocate the use of facilitators in meetings. A facilitator is an outsider with no vested interest other than to assure the best meeting possible. Facilitators express no opinion on the content issues of the meeting while at the same time keeping the meeting focused.

One of the most significant problems in removing the inhibitions to collaboration is the notion of losing one's identity to a "group mind." This problem is very real and must be addressed by the workgroup. Whole studies have been done on this phenomenon, widely known as "groupthink." Michael Schrage speaks briefly to the phenomenon: "The individual ego evolves into a collective ego. The focus is no longer on who said what but on what's on the screen."

The short answer to the best way of dealing with the groupthink phenomenon is to encourage diversity within the collaborative workgroup. Specific ways of dealing with the groupthink problem will be addressed later in Chapter 8: Problems with Groupware.

Redefining the Meeting Ecology

Now that we know the best way to create an effective environment for collaboration is to change the ecology of the environment, the question remains how to do so.

Begin by identifying the type of organization you're working within. All groups and organizations are either hierarchical or horizontal. The larger the organization, the more likely it is to be a hierarchy. Horizontal organizations are boards of directors, legislatures, and the like.

In a hierarchical organization, at any given level of the hierarchy, a single person is responsible for making final decisions.

In a horizontal organization, at any given level, a group of individuals is responsible for making decisions. Usually decisions can only be reached by consensus and quorum.

After you've established the organizational type you're working with, identify your role within the meeting. Regardless of organizational type there are only four possible roles for each meeting participant:

- Hierarchical meeting leader
- Hierarchical meeting participant
- Horizontal meeting leader
- Horizontal meeting participant

When you've identified your role within the meeting structure, it's time to begin to identify what the meeting is to accomplish. It's important to keep in mind four traditional problems most meetings are confronted with:

- Meeting participants tend to raise different issues simultaneously. Each participant is working from his or her own agenda, with a preconceived notion of priorities and potential solutions. Lots of energy is lost in the

meeting because everyone is focused on his or her own issue to the exclusion of all other issues. This effect is shown in Figure 6.1

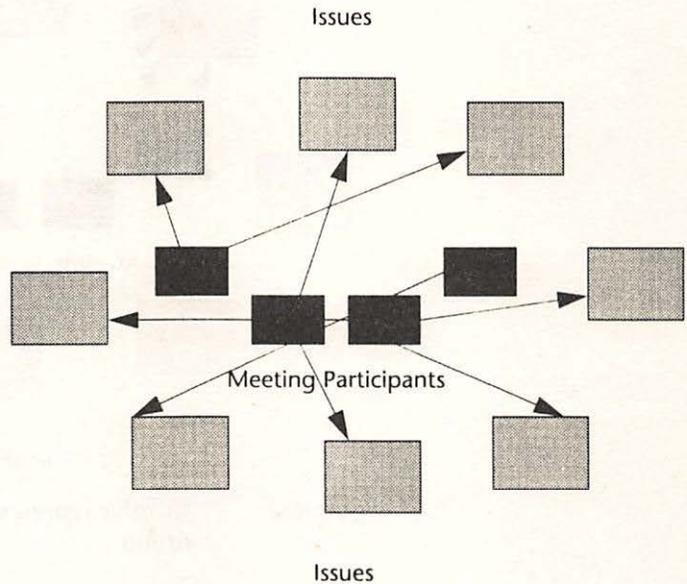


Figure 6.1 *Graphic representation of a meeting before facilitation.*

- Even when participants are really trying to help, each person can perceive the problem differently and usually has a preconceived solution to the problem or problems they perceive. The meeting will be a failure until all the participants agree to focus on one problem at a time.
- This group focus results in a level of creativity that is greater than the sum of its parts. The effects of facilitation are shown in Figure 6.2.

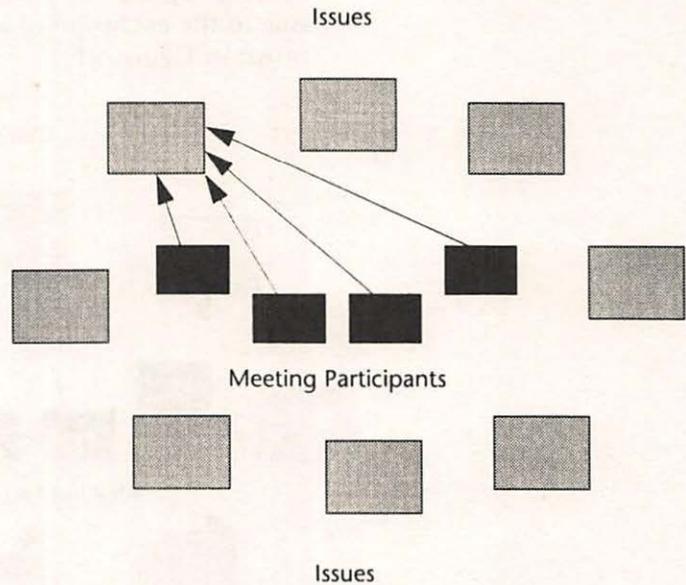


Figure 6.2 *Graphic representation of a meeting after facilitation.*

- Meeting participants tend to have difficulty differentiating the *content* of a meeting from the *process* of the meeting.
 - Focus in a meeting can only be achieved by the group agreeing on what to discuss (the content) as well as how to discuss it (the process).
- Meeting participants tend to tell the leader or the person who called the meeting what they think he or she wants to hear.
 - Managers shouldn't run their own meetings. They should turn "control" of the meeting over to a facilitator and act in a participatory rather than controlling role.

- Meeting participants usually have vested interests that may or may not be apparent. This is not a problem unless participants are attacked for their ideas.
 - People won't participate or contribute if they feel they are going to be attacked.
 - A similar problem is encountered in groups larger than about four or five people. All the participants spend a significant amount of energy trying to appropriately jump into the flow of the conversation.

These problems are effectively addressed by replacing the meeting manager with a facilitator and implementing a system of group memory. Groupware's shared space is a perfect example of the concept of a group memory.

The Role of the Facilitator

It's important to recognize that process control often results in content control. It's impossible to run a non-manipulative meeting if you have a vested interest in the subject matter. Turning "control" of the meeting over to a facilitator helps ensure that the meeting will be fair and non-manipulative.

The facilitator is responsible for maintaining a balanced conversation flow, protecting group members from personal attack, and for ensuring that each participant is heard.

The facilitator must remain neutral, not evaluating the ideas of the participants or interjecting his or her own ideas into the meeting.

The Role of Group Memory

Most of us can only hold seven pieces of information in our short-term memories at one time. There's no reason to expect the collective memory of a group to be any different.

Traditionally, meetings have had long-term memories in the form of meeting minutes. These minutes are flawed in that they represent only a single person's recollection of what took place at the meeting and they are not distributed until after the meeting is over.

Most groups don't have a collective short-term memory and the participants are spending most of their mental energy trying to hold on to their own ideas.

The Macintosh and its associated software can be used to supply a real-time group memory that is both accurate and immediate. This group memory helps provide a tangible focus point for the group members.

The key to changing the ecology of a meeting is to remove the presentation aids that are static and replace them with a shared space that is dynamic and accessible to all members of the meeting. The most appropriate shared space currently available to us is a network of interlinked computer screens.

The shared space inherent in the network of computers changes the ecology of the collaboration. The traditional meeting—discussion of a topic, reaching a conclusion, and making a decision—is replaced with shared creation. In addition, each member of the collaborative team walks away from the meeting with a tangible product—a laser printout of the work performed in the collaborative session. The immediacy and tangibility of the results reinforces creation aspect of the work. As Michael Schrage notes, "the focus shifts from planning and reviewing work to actually doing work."

The fundamental idea behind the group memory concept is to shift the attention of workgroup members from them-

selves and their agendas to the computer screen. Schrage points out that this attention shift “redefines all the relationships within the group because it [the computer screen] becomes a new medium of communication.”

Michael Schrage provides a lucid example of the facilitator employing the shared space and group mind concepts to change the ecology of a meeting. “For the moment, let’s assume that a stenographer (or a facilitator) is typing up the discussions at a meeting even as people speak, taking ‘living minutes.’ As people talk, their comments appear on the screen, the main source of light in the room. People can see themselves being heard. What appears on-screen changes constantly as people talk. These changes command attention. Everybody sees this at the same time. Because everyone has access to the screen, the screen becomes the community medium. It becomes the shared space. Shared space transforms the ecology of a meeting as surely as a change in season transforms the ecology of the rain forest. This ecology of shared space completely transforms the perceptions of conversation, information, communication, and collaboration. The perceptual bias of the medium is toward a melding of oral and visual communication.”

Arriving at Consensus

The goal of the transformed meeting ecology is to arrive at a group consensus. Consensus does not mean compromise or negotiation. Nor does it necessarily mean agreement. A consensus is reached when a solution to a problem is arrived at that does not compromise any participant’s convictions or needs. As a participant, you may not believe that the consensus is the best solution, but you feel that it’s one you can live with and support without losing anything important. Negotiation is bargaining, not collaboration.

Jay Hall discovered that a decision arrived at by a group is qualitatively better than a decision arrived at individually so long as the following five guidelines are followed.

- Avoid arguing for your own views.
- Do not assume that any decision results in a “winner” and a “loser.”
- Do not change your mind to avoid conflict or to reach an agreement.
- Avoid such cop-outs as majority voting, bargaining, or a coin-flip.
- Expect differences of opinions; even strong ones.

Benefits of the New Meeting Ecology

One of the strongest advantages of establishing a new meeting ecology through collaborative computing is that positions, facts, and assertions can be examined easily for context as well as accuracy. Bernard DeKoven, the technographer/collaboration facilitator discussed earlier, has identified three phases in the structure of computer-augmented collaborations, especially those based on brainstorming. He calls this inherent structure the “C-Cycle.”

- *Collecting comments.* Ideas, comments, questions, and concerns are listed in the group memory. Since all the comments are being recorded, the pressure of remembering everything is removed from the workgroup members. Similarly, participants don't feel like they have to hurry to get their comments out before losing their train of thought. Generally this phase lasts for about an hour, sometimes longer.
- *Connecting comments.* Workgroup members draw links, parallels, and connections between the comments that have been recorded in the group memory. Disagree-

ments invariably arise, and this is the time to address them. Arguments range from semantics—how the comments should be phrased—to how comments should be connected. During this phase, a consensual structure is formed with grouped comments forming categories.

- *Correcting comments.* This final phase in DeKoven's C-Cycle is the most difficult. In the correction phase, categories of grouped comments are ranked in order of importance, causing disagreements that are sometimes quite intense. The good news is that the disagreements that arise are real ones, tied to differences of opinion rather than misunderstandings, misperceptions, or ego conflicts.

Managing the Information Flow

Most organizations do not understand how to manage the flow of information. They are tied to an archaic model of information; believing that information is tangible and can be managed.

Information is intangible and has little value in and of itself; the value is created by the people who work with it. Information, in and of itself, can no more be managed than a tidal wave. The flow of information, however can be managed.

Irene Grief, of Lotus Development Corporation, has identified two dominant models that are used by most organizations to handle information. According to Grief, neither approach supports workgroup collaboration.

- *Personal computing approach.* The underlying concept to the personal computing approach is the encouragement of creativity and innovation.
- *Systems management approach.* The underlying concept of the systems management approach is to control the flow of data.

The design of most collaborative computing applications encourages communication but detracts from actual collaboration.

Desktop publishing tools are a prime example. While making printed communication easier and more affordable, the tools did little to enhance the collaborative nature of the publishing process. Desktop presentation and multimedia tools are following in the same tradition. These tools do little than add a glossier envelope within which to package information.

Designing tools and workspaces for the collaborative process requires a shift from models based on distribution networks to shared space networks. The issue evolves from simply processing information to creating information. Unfortunately, this cannot be accomplished by simply tacking on networking capabilities to existing software applications designed for individual use.

Michael Schrage identifies this problem as one of quality vs. quantity. "The issue is the quality of collaboration, not the quantity and frequency of communication.... People who build information tools are increasingly sensitive to their design bias and limitations. The personal computer software industry, which rose to multibillion-dollar status on the strength of enhancing personal computations and communication, is now aggressively exploring the potential of using software to support groups. This so-called groupware takes the view that personal computer software can be extended to groups; that software can be grafted onto how people work together. That may be true, but the reality is that most software designers know far more about collaborating with computers than with human beings. And, of course, they design around that core of knowledge. The result is that most of what's called 'groupware' is designed more like a prosthetic than an augmentation of how people naturally collaborate. Technology is always more likely to

succeed when it extends what people tend to do naturally rather than require a new repertoire of behaviors.”

Collaborative tools designed to augment the work of two-person work teams will likely be the area of the greatest focus for the near term. Kleiner, Perkins, Caulfield, & Byers venture capitalist John Doerr has identified this trend as something he calls “pairware.”

Michael Schrage goes even further when he writes about the nature of collaborative tools deflating the myths of information organization and management. “The rise of collaborative tools will puncture the myths of the information organization as the techno-vision of the future,” he writes. “The clichés aren’t wrong—they’re just not right. The clichés used to be ‘It’s not what you know—it’s who you know.’ Or ‘It’s not what you know or who you know but when you know it.’ The collaborative cliché will be ‘It’s not just what you know, who you know or when you know it, it’s how you create value with other people.’”

Agendas for collaboration are important. All participants should know well in advance what to expect before the collaborative session. The agenda helps settle procedural questions in advance of the session or meeting and allows meetings to be run in a much tighter manner. Meetings, therefore, become shorter and more effective.

Alvin Toffler maintains that there is an intrinsic relationship between knowledge and power in our society. He writes that the “way we organize knowledge frequently determines the way we organize people—and vice versa.”

When most of our business organizations were hierarchical, knowledge was perceived to be hierarchical as well. As our organizations now tend toward the horizontal, so does our knowledge. It’s important to be careful, then, about how you go about organizing knowledge within your organization.

Using Satellites and Independents

One of the most compelling trends as many organizations attempt to more effectively organize and manage their knowledge is the use of satellite organizations and independent contractors.

The 1980s are now recognized as a time when corporations were bent on merging with other corporations. Leveraged buyouts and hostile takeovers were the norm in this merger feeding frenzy.

A less recognized trend—the growth of the corporate profit center—also earmarked the decade. Large companies shattered into independent business units, operating as entrepreneurial small businesses. Alvin Toffler recognizes this trend as a mosaic corporate structure: “the largest corporations began shifting from monolithic internal structures to mosaics made of scores, even hundreds of independently accounted units.”

Unfortunately, most of these mosaics were nothing more than miniature versions of the original monolith. It's readily apparent that the logical next step from the mosaic corporation is the loosely-tied network of independent contractors. These independents would come together for a project, and then dissipate when their work was finished. A logical next step, yes, but it remains to be seen whether or not corporations take it.

Designing the Physical Space for Collaboration

The collaborative environment can be enhanced significantly by designing the physical space with collaborative activities in mind.

Semicircular tables or desks make the most appropriate real-time collaborative environment. The tables or desks should

be arranged in such a way as to radiate toward the shared space or group memory rather than each other. The arrangement should be such that the tables or desks face away from any doors so that people entering or leaving are not disruptive.

One of the striking differences between high school and college for me was that all high school classrooms are arranged with the desks facing the door. College classrooms are arranged with the desks facing away from the door. I'm sure that's why high school felt so much like a prison and college felt much more open. Entering or leaving a high school classroom during class is an experience in non-verbal intimidation.

Circular arrangements are appropriate for casual or informal gatherings where direct eye contact is desired. This arrangement is not conducive to problem solving, however, as it tends to heighten the mood of any group. Any negative energy will be aimed directly at one of the participants.

Good, even lighting is a necessity. The best is indirect incandescent lighting.

The room should be somewhat isolated from the rest of activity in the building and it should be as soundproof as possible.

Finally, the meeting space should be as politically neutral as possible. Where meetings are held—your office or mine—is a political matter.

Tenets for Success in the Collaborative Environment

Michael Schrage has identified 13 tenets of success in the collaborative environment. Schrage refers to these 13 points as "collaboration design themes."

- *Competence.* Successful collaboration depends on the competence of the workgroup that has been assembled. Sometimes it's surprisingly easy to overlook this basic point.
- *A shared, understood goal.* Schrage claims that collaborative activities are "classic examples of management by objective." In order for a collaboration to be effective all workgroup members must share a common goal.
- *Mutual respect, tolerance, and trust.* Schrage points out that friendship, interestingly enough, is not crucial to an effective collaboration. He points out that "successful collaborators tend to ignore the more irritating quirks and idiosyncrasies of their colleagues." What is necessary for a successful collaboration, however, is a basic respect for other workgroup members.
- *Creation and manipulation of shared spaces.* The media used for the actual collaboration—in most cases, a network of Macintoshes—becomes a shared space for the workgroup members. Effective and well-managed shared space is crucial to the success of the collaborative activity.
- *Multiple forms of representation.* The old saying that a picture is worth a thousand words is amplified in the collaborative environment. Drawings, even animations, serve as multiple forms of representation of what a workgroup member is saying. According to Schrage, "because there frequently is confusion over language, collaborators look to other representations to triangulate their perceptions and impressions." The multiple representations themselves form a web of reference points, associations, and cross-references that in itself has great value to the collaboration.
- *Playing with the representations.* Questions and problems are treated as opportunities for exploration and learning rather than stumbling blocks and barricades in the effective collaborative environment. Experimentation

is the rule of the day, and as Schrage points out, “the multiple representations and the shared spaces serve, in effect, as a conceptual and technical playground for the collaborators.”

- *Continuous but not continual communication.* Effective collaboration tends to flow like good conversation. Except in cases of dire emergency—Schrage uses the example of an airline cockpit emergency and a hospital operating room—collaboration does not require constant communication. The communicative flow that evolves within the collaborative activity “maximizes both flexibility and spontaneity—two qualities of communication that successful collaborators stress are essential,” according to Schrage.
- *Formal and informal environments.* Room and allowance must be made for both formal and informal environments in order for effective collaboration to take place. Most effective collaborators tend to agree that a sense (and place) of both spontaneity and structure are necessary for success.
- *Clear lines of responsibility but no restrictive boundaries.* Schrage states emphatically that “collaborative relationships are not managed.” Nevertheless, someone must assume eventual responsibility for the outcome of the collaboration. Because the end-result of a successful collaboration is a shared understanding and solution, all workgroup members share an overarching responsibility for the success of the collaboration while maintaining individual responsibility for their defined duties.
- *Decisions do not have to be made by consensus.* Effective and successful collaboration does not require consensus and certainly not agreement. As shown earlier, the shared space and group memory aspects of the collaborative environment help to depersonalize arguments and disagreements and focus them on the issues rather than the personalities involved.

- *Physical presence is not necessary.* Physical proximity makes effective collaboration only marginally easier. While face-to-face contact can help a collaboration, advances in technology have diluted the advantage of geographic proximity.
- *Selective use of outsiders for complementary insights and information.* Schrage cites that “successful collaborations have historically relied upon a network of outside advisors who are familiar with either the technical area, the personality of the collaborators, or both.” It’s important to realize, however, that outside help cannot be forced on the workgroup; the members must be allowed to accept outside help and information on their own terms.
- *Collaborations end.* Collaborations have as their goal and end-result a shared understanding and solution of a problem or project. When that goal is met, the collaboration ends. This should be acknowledged early on, although some collaborative groups will go on to address other problems or projects. Schrage points out that “people and their collaborations are not static, inflexible machines—they are dynamic relationships that respond to changes in both environment and expectation.”

Collaborative Problem Solving

Problem solving is an inherent part of the collaborative process. Sometimes the sole reason for the collaboration is to solve a specific problem. In either case, problem solving is, at best, educated trial and error. There are no panaceas and no solution will work in all situations. There are, however, some basic tools that can be employed as a foundation for collaborative problem solving.

Problems can be reduced to situations that you want to change. Problem solving, then, can be reduced to situation changing. The most important thing to keep in mind is that problems are not solved step-by-step, although seven basic phases can be defined.

- *Acknowledging perception of the problem.* Most of us don't perceive problems, instead we perceive symptoms. The first three phases of collaborative problem solving—acknowledging perception, defining, and analyzing—are closely related. In practice, you will likely find yourself cycling between these three initial phases. What's most important to keep in mind is that you're only attempting to get a handle on the problem during these phases, not provide solutions.
 - Before any problem can be solved collaboratively, all the collaborators must agree that there is a problem and that they intend to solve it.
 - It's important to realize that in many organizations people are afraid of problems. Problems have come to somehow represent failure in our culture. In order for the collaboration to be successful, this misperception must be addressed.
 - It's common for most of the group members to perceive a problem differently. That's okay and to be expected. Because of all these different perceptions of the problem, it can be said that any "real" problem doesn't exist at all, only each individual's perception of it. Each perception is itself quite real. It's important for each perception to be legitimized and accepted by the other members of the group.
- *Defining the problem.* The problem is defined by acknowledging what it is and what it is not. The problem definition is crucial because by its very nature determines the range of solutions. The best way of dealing with this phase is by avoiding all assumptions about

the nature of the problem. Any false assumption can limit acceptable solutions.

- *Analyzing the problem.* During the problem analysis phase, break down the problem into its component parts. Get a handle on the problem by trying to understand it. The problems we're confronted with today are significantly more complex and interconnected than in the past. Therefore, it's important to recognize the interconnectedness of any problem you are working on.
- *Generating alternative solutions.* When your group concentrates on generating alternative solutions it is crucial that any proposed solution is not evaluated. Group members must be protected from personal attack. Everyone must realize that this is a time for brainstorming and free association, not a time for criticism. That will come soon enough.
- *Evaluating and judging potential solutions.* Before the alternative solutions can be judged, the group must agree on a set of evaluation criteria.
- *Making the decision.* During the decision making phase, the collaborators should review the results of their evaluation phase and commit itself to one or more decisions. Consensus is the best result and the one that should be aimed for by the group.
- *Implementing the decision.* Once the group has decided on its decision(s), an attempt to agree on appropriate implementation is necessary. Again, consensus is the best result and the one that the group should target.

Hypermedia Support is Crucial for Successful Collaboration

Most collaborations will result in the creation of a document. This document is several orders of magnitude more useful if it is a hypermedia document. A hypermedia document, for the sake of this work, can best be represented as a

hybrid combination of text, graphics, animation, sound, and perhaps video. Central to the concept of hypermedia are the intelligent, contextual links that are used to forge meaningful trails through the information content.

The next logical step after integrating hypermedia into the collaborative environment is extending it to include simulation. And the most powerful way to provide collaborative simulation tools is within the environment of virtual reality or cyberspace.

These shared spaces within the hypermedia (and by logical extension, within virtual reality) environments will likely create their own applications and devices. Marshall McLuhan aptly observed that each new medium cannibalizes the media that came before.

These concepts are expanded in Chapter 7: Collaborative Hypertext and Hypermedia and Chapter 9: Future Groupware.

Interface and Design Issues

Many collaborators assume that collaborative tools will develop around their own contexts. Michael Schrage goes so far as to claim that “collaborative tools will give us collaborative metaphors.”

Maybe, maybe not. It's entirely possible that questionably designed first-generation collaborative tools may themselves spawn appropriate metaphors for second- and third-generation collaborative tools. But then again, maybe not. What happens if no one buys into the first-generation technology. Look how long it's taken Engelbart's tools to be appreciated and built upon. It's important to design collaborative tools effectively and with appropriate interface metaphors, especially in the earliest stages. This will not happen by accident.

Michael Schrage insists that collaborative design has yet to find its genius and implies that it will take a collaboration to implement appropriate tools and environments. This is similar in concept to the "bootstrap effect" Douglas Engelbart advocated. Schrage has identified two basic design themes that are crucial in the design of tools and environments to support collaborations.

- Collaborative tools and environments must be designed to support relationships.
- Collaborative tools and environments must be designed to support processes not products.

Collaborative environments should be designed so as to support a seamless portability of the collaboration itself. The environment should be designed specifically to support collaborations wherever and whenever they occur. Such environments would themselves be regenerative in nature and would spawn other, even more appropriate, environments that we can so far barely imagine.

Seven general questions should be addressed during the design and implementation phases of developing a collaborative environment:

- What happens during one cycle of the business or industry?
- How do information and deliverables get passed between workgroup members?
- What kinds of information gets distributed?
- What are the various jobs and who does them?
- How are projects tracked?
- Where are the bottlenecks?
- What are common job frustrations that the technology can help address?

Synchronous and asynchronous communications both must be provided for.

- *Synchronous communication* occurs when participants are connected to each other by technology or geographic proximity as in meetings or telephone calls that take place in real-time.
- *Asynchronous communication* occurs when the communication is received after they have been sent as in electronic mail and bulletin boards.

According to Harry Vertelney, of Apple's Advanced Technology Group, progress and decisions during synchronous communication are made on the basis of "availability of data, verbal communications, visual cues, and reactions of participants." Because of this, synchronous communications for the collaborative environment requires at least four channels: voice, sound, vision, and data.

When designing user interfaces for collaborative shared spaces, the designer must be aware of the fundamental tension between shared space and personal space. The interface must be designed so as to allow for quick redress of misunderstandings between collaborators.

Keep in mind that various media influence the course of a meeting because they interact strongly with participants' resources for communication and memory. Just having the communications windows and shared spaces will change the behavior of people.

In any group dynamic, various social, political, motivational, and economic factors must be addressed. Bear in mind that group processes are exceedingly difficult to study. Within any group, the members of the group take on different roles at different times.

In a group editing system, for example, various members will be writers/editors and senders/recipients at different times.

Groupware software must support several people working collaboratively on a task and, at the same time, it must also address the members' different and changing roles.

Chapter Seven

Collaborative Hypertext and Hypermedia

Because most collaborations will result in the creation of a document, it's important to take a look at the most appropriate document forms for the end-result of the collaborative process. This resulting document will be several orders of magnitude more useful if it is a hypertext or hypermedia document.

Hypertext is a way of organizing computer-based text that allows the interlinking of information. These links may be traversed rapidly by following associative trails through the information space.

Traditional text, with the exception of footnotes and cross-references, is linear. Our memory, however, is based on association; we form idiosyncratic links in our minds that are quite unpredictable. Traditional text does, however, provide the impetus for hypertext through the following familiar conventions:

- Footnotes
- Cross-references
- Post-it notes
- Library card catalogs

- Anthologies
- Tables of contents
- Indices
- Commentary
- Annotation
- Marginalia
- Quotes

These traditional text forms and conventions have computer-based equivalents:

- Pop-up notes
- Linked windows
- Outliners (stretch-text)
- Linked cards
- Simulations
- Overview maps
- Indexed branch points
- Semantic networks
- Pop-up notes and linked windows
- User modifiable fields
- Relational database aspects

Ted Nelson, who coined the term, called hypertext “nonlinear writing.” Nelson envisioned a “docuverse” that would contain all of the world’s documents interlinked in hypertext fashion.

He called his implementation of this system Xanadu. Nelson describes Xanadu as “a world-wide network, intended to serve hundreds of millions of users simultaneously from the corpus of the world’s stored writings, graphics, and data....

Xanadu is not a large centralized software system but rather an idea for software for running a decentralized network."

The Xanadu system will change the way we publish information:

- Authorship credit can be allocated automatically.
- Royalties, based on readers' use of documents, could be allocated automatically by the system.
- Any document can be quoted from or linked to while maintaining the integrity of the original.

Doug Engelbart is responsible for many aspects of both collaborative computing and hypertext. Engelbart's Augment system was the first operational collaborative hypertext environment.

Augment consisted of both shared and personal space with private mail, personal and shared files, personal and shared working space, personal and shared notes, etc. It also boasted community storage of intelligence, links, indices, reports, notes, and meeting transcripts. Several users could access and change these files simultaneously.

Documents on the Augment system could be "frozen" and archived in journal format. This process followed the model of professional journal publishing and was the first workable instance of electronic publishing and hypertext publishing.

Hypermedia is a superset of hypertext technologies and principles. Hypermedia adds graphics, video, film, animation, and sound to the basic characteristics of hypertext. In this work, the terms hypertext and hypermedia will be used interchangeably.

Hypertext is recognized by three basic characteristics:

- A network of textual and/or graphical nodes.

- Software facility and methodology for creating links between nodes.
- Interface tools that enable the user to create arbitrary linkages with buttons.

These characteristics will be discussed separately in the next three sections.

A Network of Nodes

Hypertext is made up of a network of textual and/or graphical nodes. The term *node* in the context of hypertext is different from the type of *network topology node* discussed in Chapter 3: Collaborative Computing With Groupware. In this context, a node is a bit of information. The size of the node can range from a single word to an entire manuscript. A node is not limited exclusively to text; it may contain graphics, animation, film, video, or sound as well as text. A node may also be a composite node, made up of any number of layers of other subnodes.

Software Facility for Linking Nodes

The second basic characteristic of hypertext is a built-in software facility and methodology for creating links between nodes. A *link* is the connective tissue that provides the computer-supported relationship between nodes. The information is navigated by traversing links.

A hypertext is comprised of three types of links:

- *System-supplied links*. System-supplied links are created automatically by the development environment.
 - *Command and control linkages*. HyperCard, for example, automatically provides home, next, and previous linkages.

- *Automatic tables of contents.*
- *Automatic tracking of user location.* HyperCard, for example, automatically provides a record of where the user has been with the Recent function.
- *Automatic creation of user profiles.* Many training stacks, for example, provide an automatically generated user profile and suggested sequence based on user activity.
- *Author-created links.* Author-created links are those linkages supplied by the creator of the hypertext document.
 - *Links to prerequisite knowledge.* Some stacks warn or will not allow the user to advance to certain areas until they have visited prerequisite areas.
 - *Hierarchical linkages.* Hierarchical linkages are based on classification of information.
 - *Chronological linkages.* Chronological linkages are based on links across time.
- *User-created links.* User-created links are those personalized linkages that are created by the user during the course of exploring the hypertext.
 - *Detours and shortcuts.* Knowledgeable users often create detours around background material to allow quick navigation to more specialized information.
 - *Annotation and commentary.* Notes and reminders, as well as commentary, can be included by the user and linked to the underlying material.
 - *Analogical links.* Users can provide their own analogies between various nodes within the hypertext.
 - *Insertions and additions.* The user can create completely new material within the existing hypertext.

- *External links.* Users can create external links to other knowledge bases.

All hypertext links, regardless of type—system-, author-, or user-created—share two basic characteristics: directionality and multiplicity.

Link Directionality

Link directionality refers to the direction in which the link traverses the information. A one-way, or uni-directional link takes the user in one direction but doesn't bring him back. A two-way, or bi-directional link provides the user with a way to travel to the end of a link as well as a return route.

Link Multiplicity

Especially large hypertext systems would suffer from an overabundance of links between nodes, overwhelming the user with the number of linkage choices. Some researchers have proposed a four-level system of dynamic filters to deal with this problem. Such a system would allow the user to select from a subset of available links based on preference.

- *Filter by category.* A list of categories would be provided from which the user would select a pre-defined path through the available links.
- *Filter by vote.* The system would maintain a record of the preferences of users and would present each new user with a group of selections based on the previous experience of other users. At the end the system would query the user for his or her preferences to add to the preference record.
- *Filter by level of expertise.* Subject-matter experts would provide a set of customized links through the material.

- *Filter by menu.* A list of options would be presented to the user at each junction within the material.

Buttons for Arbitrary Links

A button is a specific point in the hypertext that enable the user to move between links with the click of a mouse. Buttons are what make the links in a hypertext visible (and usable) to the user.

Buttons within the hypertext allow the user to browse rapidly through the information and to follow the author's nonlinear discourse. Hypertext authors share the heritage of past authors who have attempted nonsequential writing with more traditional tools.

Buttons used to traverse linkages within the material also enable the author to provide read-only hypertexts. This is valuable because certain materials, such as procedural manuals, training materials, and documentation need to be unalterable.

Most hypertexts, however, benefit from user-modification. The result is a body of information that is personalized and of significantly more use to the individual.

Six Modes of Hypertext

Hypertext can be used in six basic modes ranging from grazing or browsing to complete analysis and cross-referencing. A brief discussion of the six modes is provided in the bulleted paragraphs that follow.

- *Browsing or grazing.* Browsing is an activity indicated by skimming quickly across large amounts of information looking for topics of interest. Grazing is identified by the user simply wandering through and absorbing

whatever is close within his or her position within the material. Some observers have likened it to slow-motion browsing.

- *Training.* A training mode is exemplified by a achievement goals that have been specified by someone other than the user.
- *Online help.* The help mode is a specialized subset of the referencing mode that is recognized by much smaller amounts of information provided to aid the user in his or her current task.
- *Presentation.* Presentation modes are identified by a system-provided overview or summary of the current subject matter.
- *Learning and analysis.* Learning and analysis modes are used to actively solve problems, analyze information, or make decisions.
- *Referencing.* The referencing mode is used to retrieve specific information from within the hypertext. In current hypertext implementations, the referencing mode relies heavily on more traditional methods of information retrieval—indices, tables, etc.—than associative linkage of information.

In addition to the six modes of hypertext mentioned above, hypertext systems can be either private or public in nature.

Private hypertext systems are designed for use by small, designated groups. Access to the hypertext is denied to anyone outside of the group. A company's research and development knowledge database, for example, would be a private hypertext system.

Public hypertext systems are designed for use by large numbers of individuals, resulting in a significantly larger, shared hypertext system.

Collaborative Hypertext Design Issues

Collaborators, when creating a hypermedia document, must address design issues on both the development and use levels.

Development Design Considerations

Several issues must be addressed in the early stages of actually creating the hypertext. The earlier these considerations are decided upon, the smoother the document creation process will be for the collaborative workgroup.

What size should the nodes be and what should they contain? *Granularity* is the amount of information contained in a node. Popular hypertext systems sport varying levels of granularity including the following: a rough approximation of paper metaphors (index cards, paragraphs, articles, chapters, sections, and books); screens; single sentences; scrolling text; and variable-sized windows.

What sorts of information should the links connect? Should links connect only hierarchical information, or should keywords and references be included as well?

Where should buttons be placed? The user must be able to recognize a button and an overabundance of buttons will cause frustration and confusion. Some designers tend to "oversell" the buttons within their hypertext documents. The user must be made to feel that he or she can "trust" the buttons that are provided. "Is this link worth the effort?" "Will I get lost?" "Will I be provided with a way to get back?" etc. are all questions that the designer must answer during the development cycle.

Writing hypertext requires skills in addition to those of traditional writing forms. Hypertext authors must not only be

excellent writers; they must also possess (or have access to) design, graphics, interface, and programming skills.

Creation and maintenance of hypertexts is a labor-intensive process. Most links within the hypertext must be created by hand. Additionally, those links must constantly be maintained and updated as the information base changes. The integrity of the links (i.e., the link actually points to something) must be painstakingly maintained.

Version maintenance and control is difficult, and few tools exist to automate the process. Especially in workgroup situations, version maintenance is a major issue. Who decides when a new version number is required, and how are the dependent documents notified of the upgrade? Similarly, who decides whether or not to carry forward links in early draft materials?

User Design Considerations

It's imperative that the hypertext document be created with the eventual user or users in mind. Sometimes the user base is the workgroup itself, but more often the document is to be distributed externally or at least within other parts of the organization.

Normal reading cues can be inadequate in the hypertext environment. Novice users will tend to give up too soon if the document is not engaging.

Readers may be overwhelmed by the inappropriate use of pronouns within the hypertext. They may traverse a link to a statement like, "He was instrumental in the forwarding of this principle and is widely credited with the discovery of xxx...." The reader has no wayfinding aid to identify who, exactly, the pronoun "he" is referring to. The solution is to take care in identifying pronouns and to structure the infor-

mation in such a way as to be as highly self-contained as possible.

Readers may be uncomfortable with the nature of an underlying metaphor, running throughout the hypertext, that they have entered without proper foundation. This is addressed by providing an adequate foundation for any metaphor employed within any hypertext.

Some research indicates that readers build a hierarchical structure around the document in their minds. If this structure is not reinforced by the document structure itself, the reader becomes frustrated. Providing appropriate hierarchy and a system of readily identifiable headings and subheadings will alleviate this problem.

Readers may encounter inappropriate sequences in the middle of traversing a link. For example, a link may dump the reader in the middle of a sequential description of a procedure: "The third step is to...." The solution is to take care in placing sequences within the hypertext. Sequences are most effective when clearly labelled and linked to and from in their entirety.

Readers may encounter transitional structures that appear to loop back upon themselves or into transitional dead ends. For example, "the discussion of this topic is beyond the scope of this document, please refer to xxx in document yyy." This is easily addressed by placing transitional structures, such as the one described above, appropriately within the hypertext. Transitions are best handled in overviews and introductions.

Research indicates that just as people learn differently, they also read differently. Left-brained readers learn best by travelling sequentially through a document from beginning to end. They tend to be very uncomfortable with hypertext because they resent being forced to make link choices. Further, they resent the introduction and use of concepts and terminology they don't understand. This type of user tends to

benefit most from definition-type links and have problems with overviews or concept maps. Right-brained readers, on the other hand, tend to learn best by skipping around in material. They generally try to get a grasp of the overall concept before delving into details. Right-brained users tend to adapt quickly to well-designed hypertext environments.

Metacognition is the ability to learn how to learn. Poor metacognition skills are rampant within the typical business environment, and the lack of these skills can radically limit the effectiveness of hypertext in these settings. The single biggest problem people with poor metacognition skills have with hypertext is that they feel overwhelmed with what appears to be the unmanageable nature of hypertext.

There is some early evidence that indicates that this effect can be managed by increasing the granularity of the hypertext through the use of appropriately sized chunks of information.

Early hypertext research identified the “lost in hyperspace” phenomenon as the single biggest problem most novice hypertext users face. The problem is that the “lost in hyperspace” problem isn’t really a problem at all. It’s a myth, possibly related to the metacognition skills of novice users, right-brained vs. left-brained reading/learning, or simply the user’s preference. The perceived problem still persists, though, so it’s best to address it early on.

Complex hypertexts can fall victim to overchoice and cognitive overload phenomena. This is markedly different from the “lost in hyperspace” phenomenon just mentioned. Cognitive overload is a real problem and is a direct result of the information overload and information anxiety covered in Chapter 2: Driving Forces of Groupware.

The only method I’m aware of for addressing this problem is to increase the granularity of the hypertext with smaller chunks of information presented at one time.

Chapter Eight

Problems With Groupware

If the problems associated with personal computing are formidable, the problems become magnified by an order of magnitude in interpersonal computing, if only because of all the new variables that are added.

There are, of course, all kinds of problems that must be addressed when the success of a collaboration relies on the smooth operation of the technology involved. These are generally of a technological nature and are remedied through new technologies or new applications of existing technologies. Most of these sorts of problems have been covered in the earlier chapters of this book.

There are also potential interpersonal problems between the workgroup members themselves that must be resolved in order for a successful collaboration to take place. Closely tied to this set of problems is a subset of wider societal impacts that must be taken into consideration.

The problems associated with groupware, collaboration, and collaborative computing, then, can be broken down into Three general categories:

- Technological problems associated with personal and interpersonal computing.
- Potential interpersonal conflicts.
- Greater societal issues.

The two latter categories—potential interpersonal conflicts and the greater societal issues—will be addressed, in turn, in the following sections.

Potential Interpersonal Conflicts

Inarticulate workgroup members are at a distinct disadvantage to their more articulate counterparts. People who have a difficult time finding the “right words” or who have problems speaking in groups are penalized in the collaborative environment. Asynchronous meetings and communications—those carried out over a period of time, rather than in real-time—help substantially in these cases. People have time to formulate and reword their thoughts and the advantage of articulation is leveled.

Closely related to the problem of different levels of articulation within the workgroup is the tendency of some workgroup members to dominate the communications channels. This is more of a problem in the collaborative environment but is easily handled by delegating resources like screen real estate and network bandwidth on an individual basis.

Similarly, some workgroup members tend to view shared space as private territory. These individuals are nothing less than egomaniacs and must be kept on a short leash. The shared space—and all the resources that comprise it—must be seen as community property.

Workload vs. Perceived Benefit

There is a genuine and perceived disparity between who does the work and who gets the benefit in many groupware scenarios. Most groupware applications require additional work from some users who never see benefit from their efforts.

Consider the automatic meeting scheduling feature popular in most current groupware applications. The benefit is heavily weighted in favor of the person setting up the meeting. He or she selects a time and creates a distribution list for other attendees who must maintain individual schedule files which must be made available on the network. These electronic calendars are used mainly as a mandated communication from managers rather than a truly collaborative tool.

Sometimes groupware technology violates the corporate ethic or fabric of the organization. An automatic scheduling system, for example, that automatically schedules a meeting during what appears to be a manager's idle time, will quickly be rejected by the manager.

Adding voice annotation to a document benefits the speaker of the message and requires additional work on the part of the recipient. Speech is faster to produce for the speaker, but harder to understand than typed material for the recipient. In addition, voice annotations cannot be scanned or reviewed easily and are generally harder to manipulate.

Provision also needs to be made for error and exception handling as well as allowing room for improvisation. Humans make errors, handle exceptions, and improvise quite effectively on their own, but less so in a group situation. Groupware tools have to make allowance for this aspect of collaboration.

Electronic mail, file service, and print service can safely be called the most successful of the currently implemented

groupware technologies. Evidence would indicate that this is so because these implementations provide an equitable balance between who does the work and who reaps the benefits.

Electronic mail in particular, because it generally follows a conversational format, is compatible with existing social conventions. The exceptions to this are clearly identified as flaming and junk e-mail. Another problem exists in some organizations with the ability of an individual to send a message to anyone and everyone within the organization.

Because electronic mail, file service, and print service are asynchronous activities, the exception handling capabilities are adequate.

Problems with Evaluating Tools

We generally haven't yet learned how to effectively evaluate tools designed to enhance the collaborative process. Even in the best cases we tend to use the wrong yardstick.

It's exceptionally hard to create an evaluation setting that accurately reflects the social, economic, political, and motivational aspects that are inherent in a group dynamic.

It may be best to use methods associated with psychology, sociology, and anthropology to evaluate groupware.

Hidden Agendas and the Gatekeeper Effect

Hidden agendas must not be tolerated and must be confronted head on. Thankfully, most of this is handled by the inherent nature of the collaborative computing medium itself. The shared space must be kept as politically neutral as possible.

Michael Schrage offers an excellent assessment of this potential conflict: "Political participants and group members with hidden agendas find the new environment distasteful. By decoupling contributions from personalities, charismatic and political power is diffused. Politics and gamesmanship never vanish, of course, but unless the player wants to pull a coup d'état and capture the screen, it's very difficult to schmooze, flatter, or intimidate someone into going along with a point of view. A player with a hidden agenda also quickly discovers that since the purpose of the process is to flush out what people think, his agenda isn't likely to stay hidden for long. Worse yet, it becomes transparently obvious what that player is trying to do as that hidden agenda gradually leaks out on-screen."

Some workgroup members, especially those with hidden agendas, will attempt to exert gatekeeperitis. They will try to control what gets recorded and what gets discarded. When any workgroup member begins acting like a gatekeeper, the shared space is no longer shared and loses its effectiveness. By extension, if this is allowed to go on, the collaborative workgroup itself will quickly lose *its* effectiveness.

The single biggest interpersonal problem facing the collaborative workgroup is that the group leader will begin to act like a gatekeeper. If this happens, the collaboration dissolves into a situation where the leader imposes his thoughts and positions onto the shared space. This is precisely the reason why the best collaborative teams are formed by those leaders who don't run their own collaborations, and instead turn control of the process over to a facilitator.

Sacrificing Individuality to the Group Mind

Shifting the static and linear nature of traditional meetings to a shared space conducive to collaboration results in a synergistic merging of views and positions. This can be seen as losing one's identity to a "group mind." People want to feel

like they've contributed something substantial, something that's theirs, not just a part of the whole

It's natural for people to tend to "marry" their ideas and resist letting loose of them. Usually it's easy to manage this resistance, especially as the collaboration progresses and the workgroup members realize that they're not losing their individuality or parts of themselves. This is a natural tendency and must be addressed with understanding and patience.

Perhaps the hardest part of using any collaborative environment is assembling the right group of people. It's extremely difficult to convey nuance, body language, and facial expressions in the collaborative computing environment, and it's imperative for the workgroup members to respect each other. The workgroup gets off to a bad start if they are all thrown together for any reason other than coming together naturally. The best approach is to allow the workgroup to assemble and evolve itself based on its own needs.

Workgroup Homogeneity and The Groupthink Phenomenon

Workgroups, especially those that remain intact over a long period of time, tend to be homogeneous. This group homogeneity can lead to ineffective collaboration. Workgroups can discourage dissent without being aware they are doing so.

Irving Janis has identified this tendency of group members to think alike as "groupthink." Janis has reported nine earmarks for identifying the groupthink phenomenon.

Groups tend to form a sense of collective self-importance that is detrimental to problem solving activities. They tend to get caught up in the notion that they're on a "mission from god" and that because they're the "right" people, they will naturally make the "right" decisions. As group members rise in the organization's hierarchy, they become ever more

insulated from the realities of the lower echelons. Further, they become isolated from “outside” realities.

As a result, overly homogeneous groups tend to devolve into “us against them” mentalities. As groups become more homogeneous, they tend to develop a common set of beliefs and value structure. The groups become too comfortable; participants don’t want to disturb the feeling of solidarity with disagreement.

One of the best indications of a workgroup that has become too homogeneous is silence on the part of some workgroup members. Groups tend to take silence as agreement when in fact the silence on the part of a collaborator is disagreement.

Groupthink develops the fastest when the group is under pressure to make decisions. Once a groupthink decision has been made, there tends to be an unspoken acknowledgment not to challenge the decision, even if that decision is obviously not working.

The best way of dealing with the groupthink phenomenon is to encourage diversity within the collaborative group. Workgroups themselves should be evolving entities. There’s no reason to maintain the group after its effectiveness has diminished.

Broader Social Implications

Many of the general societal problems posed by groupware technologies can be related directly to social control issues which are themselves problems of understanding. We don’t fully understand the media that inundate our lives and by extension, many of us are dangerously ignorant of their impact and influence.

Social Control Issues

Harold Innis, a political economist who was more widely known and recognized as a media philosopher and harbinger of Marshall McLuhan, redefined history as a sequence of media reimagining our collective perception of time and space.

Innis replaces the notion of economic monopoly with that of information monopoly and sees the control of media by the information gatekeepers as a method of social and political control. Each media monopoly, according to Innis, is successively diluted by new media which is in turn diluted by an even newer medium. Print diluted oral traditions, radio diluted print, television diluted radio, interactive media will dilute television.

Different media, according to Innis, offer different levels and kinds of control. A medium that tends toward scarcity or that requires specialized encoding/decoding—such as high-definition television, digital video, etc.—is likely to be controlled by a small elite that can afford access to it. On the other hand, mediums such as the desktop publishing systems that have become so widespread, tend to democratize the culture by being more accessible to individuals.

The accessibility factor of a medium can shift over a period of time. As an example, consider the changes in television brought by the camcorder, satellite transmissions, and cable.

Innis also maintains that each medium has an inherent bias toward either longevity or accessibility across vast distances. He further claimed that the bias of any society's predominant medium affects the stability of the culture. A time-biased medium, such as a large centralized database, leads to a relatively small and stable community. A space-biased medium, such as the distributed virtual communities of the telecommunications networks, leads to a relatively wide and varied (and by extension less stable) community.

Marshall McLuhan expands upon the work of Innis and argues that all media are extensions of our senses. McLuhan divides history into three major ages: oral, print, and electronic.

The oral age was a closed society earmarked by a high interdependence that tended to stifle individuality. Print, according to McLuhan displaced the closed oral media monopoly and instead of relying solely on the ear, shifted the dependency to the visual. This weaning of the dependency on the oral allowed for introspection and critical thought, resulting in individuality. McLuhan saw the development of abstract thought as a direct result of shifting the emphasis from the ear (oral) to the eye (visual).

Collaborative workgroup members—especially administrators—must be aware of these tendencies and, if necessary, modify their actions accordingly.

Marshall McLuhan, in *Understanding Media*, writes that “our conventional response to all media, namely that it is how they are used that counts, is the numb stance of the technological idiot. For the ‘content’ of a medium is like the juicy piece of meat carried by the burglar to distract the watchdog of the mind.”

Our existing tools for managing information, including the most advanced collaborative computing tools (as well as those on the horizon), make it too difficult for people to noodle—to think out loud. In the truly collaborative computing environment, or in a computer-augmented meeting, all comments are recorded—even the noodles and thinking out loud.

The only ways we currently have of effectively recording the collaborative session is by manual labor or kludging together an inefficient solution based on the technology at hand. Both solutions are intrusive and anything but transparent; they require the workgroup members to constantly maintain an awareness of recording their own activities.

The solution to this problem is to meet complexity with simplicity and generality rather than the normal method of meeting complexity with complexity. We need more generalists, not more specialists.

The Macintosh needs a built-in recording mechanism that maintains a visual and auditory transcript of everything that takes place within the collaborative environment. Such a feature would have to be simple to use and it would have to operate transparently.

The foundation for such a tool is already available. Farallon Computing's MediaTracks allows you to record, edit, and annotate all activity that takes place on the Macintosh screen. Unfortunately, MediaTracks can only record a session within the memory constraints of the Macintosh doing the recording. It would have to be modified to occasionally flush out its memory to a disk file.

The best workaround currently available also employs using MediaTracks in combination with a traditional meeting recorder. The recorder would be responsible for using MediaTracks to record all appropriate screen activity and assembling an edited version of each session for the workgroup. This would be a labor- and time-intensive process, but would be well worth the effort in certain situations.

What Good are Computers, Anyway?

Everyone can tell a horror story about the "computer." An electronic funds transfer that didn't arrive on time because the "computer" went down. An automated teller machine that failed to credit an after-hours deposit because the "computer" glitched. A credit card purchase that is denied because the "computer" failed to update the payment made last month.

Everyone has a story; here's mine: I've never yet been able to finish a major project without some sort of major hardware failure. Either a hard disk crashes, or a backup fails, or a power supply goes south. Something always happens. My way of dealing with the inevitability of a major hardware failure at the worst possible time is to factor it in to my schedule.

So, what good are these computers, anyway, and why in the world should we want to entrust our most valuable information to them. And what about the inherent dangers in an information repository as vast as something like Xanadu; what about all the horror stories we hear about big brother and his even bigger computer. Or, even worse, his even bigger computer *network*. Why in the world should we choose to create even more connections between information bases? To look for an answer, let's start with looking at why, perhaps, we shouldn't.

More than 150 million individual credit records are contained in computers belonging to the five largest credit reporting companies in the United States. In the mid-1970s it was discovered that one of the largest credit reporting companies had access to the patient records of 90 percent of the nation's hospitals. More than 250,000 times each day, business subscribers of these credit reporting companies make an inquiry about the credit-worthiness of a customer and receive a response within three seconds. Every year, more than 350,000 formal complaints are registered with the single largest credit reporting company concerning the accuracy of the credit records held in its computer database. Each year, more than 100,000 of these complaints result in the information being changed.

David Burnham, in his landmark *The Rise of the Computer State*, quotes Kent Greenwalt, a professor at Columbia University's Law School, about the effect of such vast amounts of information in the hands of a few. "If there is increased surveillance and disclosure and it is not offset by greater tolerance, the casualties of modern society are likely to increase

as fewer misfits and past wrongdoers are able to find jobs and fruitful associations. The knowledge that one cannot discard one's past, that advancement in society depends heavily on a good record, will create considerable pressure for conformist actions. Many people will try harder than they do now to keep their records clean, avoid controversial or 'deviant' actions, whatever their private views and inclinations. Diversity and social vitality is almost certain to suffer, and in the long run independent private thoughts will be reduced."

As we attempt to connect computers to provide very real benefits to the society at large, we are also providing kindling and a strong draft to the already raging wildfire of what has amounted to nothing less than computer surveillance.

So, what good are computers? Computers have empowered individuals with immense personal power; the coming years promise to do the same for collaborative workgroups. Computers have returned philosophy to an integrated part of everyday life. When we work with computers we are forced to ask questions about the world around us:

- What is thought?
- What is intelligence?
- What is the mind?
- What is reality?

Computers are doing the same thing for our minds that the automobile did for our bodies. The automobile provided us with an extension to our physical capabilities; the computer provides us with an extension to our mental capabilities.

Most important of all, though, computers provide us with gateways to other communities. The computer allows us to create worlds that we are compelled to share with others. This marks the emergence of all sorts of things ranging from

groupware and interpersonal computing, to teleconferencing, to virtual reality.

The Mass Production and Merchandising of Information

We have been hearing from the pop-futurists about the information explosion for the better part of two decades. We have been told that we must steel ourselves against this impending blast of information. Invariably the solutions proposed by these charlatans revolve around the mass production and mass merchandizing of information.

California State University history professor, Theodore Roszak, minces few words in his description of these would-be futurologists and their musings in his 1986 book, *The Cult of Information*: "an ungainly hybrid of potted social science, Sunday supplement journalism, and soothsaying... pitched at about the intellectual level of advertising copy."

John Naisbitt's central theme in his immensely popular *Megatrends* is that we have moved from an industrial society to an information society. Familiar words, but Naisbitt's view of the information society is markedly different from the ones offered by the likes of Stanley Davis, Alvin Toffler, and Richard Saul Wurman. Naisbitt maintains that "we now mass-produce information the way we used to mass-produce cars. In the information society, we have systematized the production of knowledge and amplified our brain power. To use an industrial metaphor, we now mass-produce knowledge and this knowledge is the driving force of our economy."

Here Naisbitt has clearly made the jump from the Sunday supplement into the realm of intentional obfuscation. While the society is deeply involved in the process of mass producing information, it's absurd to make the leap to the mass production of knowledge. "Information, even when it moves at the speed of light," notes Roszak, "is no more than

it has ever been: discrete little bundles of fact, sometimes useful, sometimes trivial, and never the substance of thought.”

There are clear distinctions to be made between the notion of information and knowledge. According to Roszak, we’ve always judged the value of knowledge on three factors:

- depth
- originality
- excellence

With the inflated valuation of information comes a blatant, if unconscious, attempt to blur the distinction between what constitutes knowledge and what is merely information.

Access Issues

The economic impact of high technology advancements—such as those inherent in groupware and collaborative computing—tends to exacerbate the growing chasm in society between the haves and have-nots. As we quickly approach the danger of becoming a two-tiered society, high technology enterprises in general, and the computer industry in particular, show no overt signs of attempting to check the slide.

Everett Rogers and Judith Larsen, in *Silicon Valley Fever: Growth of High Technology Culture*, point out that “Silicon Valley means low-wage, deadend jobs, unskilled, tedious work, and exposure to some of the most dangerous occupational health hazards in all of American industry. It is a dark side to the sparkling laboratories that neither barbecues, balloons, nor paid sabbaticals can hide.”

Alvin Toffler maintains that access (and privacy, for that matter) is a political issue. “Designing a system so that it serves one department better than another is a political act.

Even timing is political, if one unit gets a lower communications priority than another, so that it must wait for service. The allocation of cost is always a power issue. ... Should companies establish internal 'information councils' or even 'legislatures,' to write the laws governing information rights, responsibilities, and access? Should unions share in this decision-making? Do we need 'corporate courts' to settle disputes over security and access? Do we need 'information ethicists' to define a new informational morality? ... Will we eventually need an explicit Bill of Electronic Information Rights?"

Information is dependent upon knowledge, not vice versa. Many in the computer industry, however, would have us believe that they are on the verge of creating intelligence in their machines. At any given time in the last two decades we have been told that artificial intelligence—and, specifically, a thinking machine capable of more and better thought than humans—was five years away, at most a decade. This is illustrated by an on-going process of what Roszak refers to as the "anthropomorphizing of the computer as a surrogate human intelligence."

As a culture we have blindly accepted and even adopted this anthropomorphism with little tension. Many Macintosh users have named their computers and spend an inordinate amount of time with their electronic friends. We have personalized our electronic workspaces with customized background screens and specialized beeps for various system functions. We speak of interfacing with our associates, and we solicit feedback from them. Most likely, this started when we began to refer to the computer as having a memory.

The anthropomorphism of the computer seems to have peaked with Apple Computer chief executive John Sculley's much-publicized Knowledge Navigator demonstration. Sculley's vision focuses on the computer as an electronic buddy with very human-like traits that carries out our bidding for us. There seems to be a glimmer of hope within the

general computer community however. In the early winter of 1988, the Society of Computer Professionals for Social Responsibility took Sculley to task for his attempts at promoting this questionable activity.

The bulk of the computer and telecommunications industries, however—the infrastructure that provides the foundation for groupware and collaborative computing technologies—seems intent on maximizing profits and focusing on increased performance. All of this, of course, is done in exchange for allowing more people access to the media.

Most telecommunications networks, for example, levy connect-time charges in single minute increments. The charges range from a low of \$5 per hour to well over \$100 per hour of connect time. This obviously takes access out of the hands of the many and places it firmly in the grasp of the few. “There is obviously a significant political public for whom the connect-time charges, let alone the price of the basic equipment, would be prohibitive. Networking may for some time to come remain a strictly middle-class medium,” writes Theodore Roszak.

Consider the activity of electronic publishing. In my own case, I publish *The Arts & Farcas Review* electronically and distribute it as shareware via several of the international telecommunications networks. For two years, the publication was distributed on floppy disk as well, and subscribers purchased a tangible product from me (the floppy disk) which I delivered to them. In a product-based economy, I would have depleted my inventory by distributing the disks. But the nature of electronic publishing is that I still had the product—in its original form—after it was distributed. In fact, as a collective body of information, it was worth more to me than it was to the subscriber.

Stewart Brand, in *The Media Lab: Inventing the Future at MIT*, addresses this paradox in a unique way. “Information wants to be free because it has become so cheap to distribute, copy,

and recombine—too cheap to meter. It wants to be expensive because it can be immeasurably valuable to the recipient. That tension will not go away.” As an example, Brand cites the Lexis online database for lawyers. The information contained in the Lexis database—in the form of judicial decisions—is in the public domain. “Multimillion dollar Lexis succeeds by charging handsomely for them,” writes Brand, “by owning copyright on the *page breaks*.”

Ithiel de Sola Pool, in his 1983 *Technologies of Freedom: On Free Speech in an Electronic Age*, takes the point of copyright one step further by declaring, “The recognition of copyright and the paying of royalties emerged from the printing press. With the arrival of electronic reproduction, these practices became unworkable. Electronic publishing is analogous not so much to the print shop of the eighteenth century as to word-of-mouth communication, to which copyright was never applied.”

Brand states emphatically what I and other electronic publishers have experienced first hand: people are reluctant to pay for quality information. According to Brand, this is because the “valuing is retroactive.” He points out, however, that people will “pay for quality of *source*, because the constancy (reliability) of source makes value somewhat predictable.”

In the United States we have an abundance of communications media open to us. Oh, we’ve had our collective bouts with the gatekeepers and the ticket-takers of the information byways, but for the most part we’ve had pretty easy sailing.

The United States government continues, for example, to insist that the radio frequencies be controlled because of a perceived scarcity. Stewart Brand offers a telling perspective on this misperception: “In Washington, D.C., there are twenty-three radio stations (licensed and regulated), eleven television stations (licensed and regulated)... and two daily

newspapers, both of them fully protected by the First Amendment. Which is the scarce medium?"

And then there was the time in 1986 when then-President Reagan decided that a lot of publicly available information living in electronic databases should be reclassified as "sensitive" and that access to the information should be restricted.

But mostly we've had pretty easy sailing.

The underlying current in all of this is rather obvious. That's probably why it's not generally perceived. The right of access to information is becoming very much a political issue. And it's not an issue, as some politicians just to the right of John Birch would have us believe, of only economics.

Stewart Brand cites Albert Bressand, the French economist, as saying that a world of information shock is coming and will crest in the 1990s. It will be similar, according to Bressand, to the oil shock of the 1970s and the world banking shock of the 1980s. It will happen "when the information providers decide to revalue what they produce."

Closely tied with the economic issues encompassing the access to information is the lowest common denominator factor. Cable television was supposed to deliver us all from the wasteland of broadcast television by providing a broader spectrum of available programming. What we got instead, with a few notable exceptions, were clones of broadcast successes, or, repackaged versions of borderline material discarded by the big three broadcast entities. Information services as well, cater to the lowest common denominator; that point at which there is not so much the most interest as the least friction.

Privacy and Security Issues

In any network situation, privacy issues will be a major concern. Every user connected to a network leaves a tell-tale crumb-trail leading into and out of the documents he or she has accessed. New trails will be forged that connect seemingly disparate bodies of information. Who will have access to this trail information? What limits will be placed on its use? What about data security?

When 1984 passed with relatively minor whimpers of Big Brotherishness from the fringes of the community-at-large, most of us felt that issues of governmental invasion of privacy and covert surveillance on the American citizenry were a thing of the past. Most of us filed those paranoid years of the 1960s and early 1970s away in a dusty corner of our gray matter and breathed an extended sigh of relief to have made it through.

We were premature in letting our guard down. The Federal Bureau of Investigation has, for at least the past 10 years, taken to recruiting librarians as little more than spies to keep an eye on what the Feds define as suspicious activity on the part of library patrons. Specifically, those who use the photocopiers and speak with foreign accents. The FBI, under its Library Awareness Program, insists that crucial national secrets are leaking out of the country and into the hands of The Enemy. The FBI, while offering no substantiation of these claims, says that the program is successful and will be continued.

When pressed, the FBI representatives point to the arrest of a Soviet United Nations employee several years ago for receiving classified information from a foreign exchange student that had been hired to photocopy material from libraries. Most librarians are quick to point out, however, that classified material is not available at any public library they know of, and that the FBI's Library Awareness Program

amounts to nothing more than invasion of privacy and maybe even First Amendment infringement.

In 1986, the FBI teamed up with the CIA and the Air Force to coerce Mead Corp. into restricting access to its Mead Data Central, the division that publishes the Nexis online database of technical publications. Mead refused, taking the position that all of the Nexis information had been previously published. The government's position, according to Jerry Yung, Mead's vice president of government relations, was that although the material had been previously published and available publicly, "once aggregated with powerful computer software, it could be used against national interests."

While Mead refused the government's request to restrict access to the Nexis database, the database that most likely set off the red flag, The National Technical Information Service, was dropped later that year.

As networks and information services become more and more inter-linked, as will doubtlessly happen since Judge Green's decision finding that the Regional Bell Operating Companies can provide information gateways, the level of data integration available to the average workgroup and individual telecommunicator will begin to grow by orders of magnitude. With that will come more and more government irritability, anxiety, and concern. When that access to information is paired with powerful filtering software and advanced hypermedia engines that are capable of semi-automatic and semi-intelligent information links, the Information Age will finally be a reality.

Access to information is a right, not a privilege as the government would have us believe. What we read and what online databases we search are our own concerns, and no one else's. This right must not be compromised in any way, and must be defended at all costs.

Electronic communications are currently not given the same protections as even first class paper mail. Most of us

continue to suppose that we have freedoms, rights, and protections that simply don't exist.

Wiretapping, for example, used to have to be done physically. There had to be a listening post somewhere nearby the person being surveilled, or at least a tape recorder in the immediate vicinity.

My friends that helped design the telephone companies' ESS switching system tell me that this is no longer the case. Now your conversations can easily be listened to without your consent or knowledge by simply sending the proper signal tones down the telephone line. From anywhere.

Computer communications are similarly vulnerable to interception. All computers, except for those specially modified, emit radio frequency waves that can be quite easily intercepted.

Finally, there is no traditional way of demonstrating the authenticity of an electronic document. There are no typographic elements or watermarks as with paper publishing.

Many believe data encryption is the most effective way of ensuring privacy and data integrity. Not so according to Ted Nelson who asserts that the Data Encryption Standard (DES) endorsed and approved by the federal government has been shown by researchers to be defeatable by the National Security Agency.

Nelson goes on to note that a better system developed at MIT was subsequently suppressed by the federal government.

Utilities such as Status*Mac and GraceLAN allow network administrators to poll workgroup members' computers and obtain in-depth inventories and profiles of each workstation. As programs like this become more widely used, conflicts will arise over an individual's right to privacy.

Similarly, corporations that promise to protect the privacy of office electronic mail are coming under fire when that privacy is compromised.

It's a crime in only three states—California, Colorado, and Florida—for an individual or a company to eavesdrop or record confidential communication without the knowledge and consent of both parties to the communication. Federal law requires only the knowledge and consent of one of the parties.

Intellectual Property Issues

As collaborative workgroups begin to publish their materials, and links to existing materials are forged, our established notions of copyright become burdensome and ineffective. Tied to the issue of copyright, intellectual property issues also have to be addressed.

Hypermedia and other forms of electronic publishing will very soon become viable alternatives for some forms of paper publishing. Surely problems involved with the intellectual property aspects of publishing will crop up and have to be addressed. Additionally, there will have to be differentiations formed between public and private publishing.

The Alexandria Institute is at the forefront of this issue and intends to capture all the world's recorded knowledge on CD-ROM and make it accessible to anyone who wants it. The future, according to the Alexandria Institute, will consist of scholars and researchers having the best libraries in the world at their instant disposal, wherever they are.

The future envisioned by the Alexandria Institute is built on a past. In this case, the Alexandrian Library created by Ptolemy in 300 B.C. serves as the foundation. Historians estimate the Alexandrian Library contained more than 500,000

papyrus rolls which were constantly revised and edited by the Alexandrian librarians. Scholars travelled great distances to make copies of the works housed in the Alexandrian Library and, in this manner—copies of copies—the manuscript collections of Europe were established. This process allowed a greater number of people to have access to the vast body of information.

Robert Kerr, the chief executive of the Alexandria Institute, is designing a copyright system that is based on a pay-for-use model. Kerr maintains that “a copyright system based on the publication, distribution, and compensation for a printed copy of a book may not be a viable model for the protection of intellectual property rights when the ‘copy’ is a single disc with the full text of 500 books representing more than 1,000 different proprietary interests.”

As illustrated earlier in this chapter, Mead Data has addressed this issue in a round-about way by declaring its copyright on the page breaks contained in its Lexis legal database. It’s also evident in Ithiel de Sola Pool’s comparison of electronic publishing with forms of word-of-mouth communications which have never qualified for intellectual property protection of any kind.

The pay-for-use model envisioned by Kerr for electronic publishing copyrights would enable information providers and publishers to deliver their products to customers for the cost of the media on which it was contained. The information provider would be compensated by an accrual of payments based on a fee structure for the actual use of the information.

Kerr maintains that such a system would benefit the information provider directly by opening up new avenues of compensation for additional uses of an existing work; allowing price discrimination among types of users; adding incentives to increase the number of users; and allowing for a low-cost inventory for publishers. Similarly, libraries would be able to maintain a significantly larger collection at a

greatly lowered cost, and individuals would enjoy access to a broader knowledge base at a more reasonable cost.

Such a system would surely be met with resistance by both publishers (who would potentially lose duplication control) and libraries themselves, who now generally pay a fixed yearly subscription rate for material delivered on CD-ROM.

Chapter Nine

Future Groupware

In 1984, Apple Computer introduced its Macintosh computer; a computer designed to be as easy to use as an appliance. Apple stressed in its marketing campaign that Macintosh was the computer for the rest of us. A large part of Apple's philosophy revolved around the notion of one computer for every person within an organization. Every Macintosh that has ever been sold comes with built-in networking capability—the foundation of groupware and collaborative computing. Apple's marketing focus has subtly shifted from personal computing to interpersonal computing.

The 1990s will see the widespread development and integration of groupware; something that many pundits were sure would occur in the 1980s. Steve Jobs has gone so far as to identify the 1980s as the age of personal computing while the 1990s will be the age of interpersonal computing. According to *Fortune* magazine, "Software that will enable people to collaborate across barriers of time and space is one of today's hottest frontiers of computer research."

Because every Macintosh sold is equipped for local area network communications, Apple's Macintosh will be at the forefront of groupware, collaborative computing, and interpersonal computing. It's likely that Apple will revive its mar-

keting prowess and attempt to lay claim to the concepts of groupware as it did four years ago with hypermedia. Apple has done a lot to further the notion of groupware and collaborative computing, but the concepts date back to the 1960s. Without doubt the first broadly successful groupware applications will appear on the Macintosh.

Virtual Reality

As noted earlier in this book, the next logical step after integrating hypermedia into the collaborative environment is extending it out to include simulation. The most powerful way to provide collaborative simulation tools is within the environment of virtual reality or cyberspace.

The shared spaces originated within the collaborative environment and extended to the virtual reality environments will likely create their own applications and devices. Marshall McLuhan aptly observed that each new medium cannibalizes the media that came before.

William Gibson is widely credited with popularizing the terms cyberspace and virtual reality in his novels. "Cyberspace. A consensual hallucination experienced daily by billions of legitimate operators, in every nation.... A graphic representation of data abstracted from the banks of every computer in the human system. Unthinkable complexity. Lines of light ranged in the nonspace of the mind, clusters, and constellations of data. Like city lights, receding...." I like John Perry Barlow's definition better: "The place you are when you're on the telephone."

NASA is deeply committed to virtual reality research to enable something called "telepresence." Telepresence is the ability to project human judgement and action into a robot located in a hostile environment.

Cyberspace technologies allow us to simulate entire worlds within the computer. In the near-term we'll access these worlds through gloves and goggles wired into the computer. Many believe that eventually these artificial props will be abandoned in favor of a direct connection between the human brain and the computer.

The experience of cyberspace is not one of a three-dimensional sensorama, but rather an immersion in another world. Aside from the science-fiction-like seductiveness of cyberspace, there are real-world applications for this next advancement in computer interfaces:

- *Scientific visualization.* The step from conventional computer modeling systems to a three-dimensional environment that is manipulable will lead to breakthroughs we can only imagine. Imagine being able to enter a simulation of a world at the molecular, or even atomic level.
- *Architecture.* Designers and architects will be able to truly visualize their designs, at true scale.
- *Telepresence.* The telepresence applications pioneered at NASA and elsewhere will lead to all sorts of real-world breakthroughs. Imagine being able to remotely control robots to rescue mine cave-in victims and explore the ocean bottoms.
- *Education.* New education techniques will be developed putting these new technologies to their highest and best use. Educators will no longer be restricted to the physical classroom. The entire concept of a field trip will be redefined as students are enabled to explore entire universes.

The potential offered by virtual reality with regard to collaboration is mind-boggling. Imagine being able to actually enter an information environment and manipulate elements as easily as we do everyday objects like a pencil or a drinking

glass. Imagine being able to pick up a piece of information and examine it from every possible angle.

These new technologies will help us all manage information more effectively. More importantly, these advanced technologies will help us learn more about ourselves and the worlds around us.

Appendix A

Bibliography

ACM, ed. *Hypertext on Hypertext*, electronic ed., 1988.

ACM, ed. *Proceedings on the Conference on Computer-Supported Cooperative Work*, December 3-5, 1986.

Adams, James. *The Care and Feeding of Ideas: A Guide to Encouraging Creativity*. Reading, Massachusetts: Addison-Wesley, 1986.

Ambros, Sueann and Kristina Hooper, eds. *Interactive Multimedia: Visions of Multimedia for Developers, Educators, & Information Providers*. Redmond, Washington: Microsoft Press, 1988.

Ambros, Sueann and Kristina Hooper, eds. *Learning with Interactive Multimedia: Developing and Using Multimedia Tools in Education*. Redmond, Washington: Microsoft Press, 1990.

Anderson, John R. *Cognitive Psychology and Its Implications*. New York: W.H. Freeman and Company, 1980.

Apple Computer Inc. *Apple Human Interface Guidelines: The Apple Desktop Interface*. Reading, Massachusetts: Addison-Wesley, 1987.

- Bair, James. "Communications in the Office of the Future: Where the Real Payoff May Be." *Business Communications Review*, Volume 9, Number 1, January-February 1979.
- Bandura, A. *Social Learning Theory*. Englewood Cliffs, New Jersey: Prentice-Hall, 1977.
- Bar, Avron. "AI: Cognition as Computation." In *The Study of Information*. ed. Fritz Machlup and Una Mansfield. New York: Wiley, 1983.
- Barrett, Edward, ed. *Text, ConText, and HyperText: Writing with and for the Computer*. Cambridge, Massachusetts: The MIT Press, 1988.
- Barrett, Edward, ed. *The Society of Text: Hypertext, Hypermedia, and the Social Construction of Information*. Cambridge, Massachusetts: The MIT Press, 1989.
- Bateson, Gregory. *Mind and Nature*. New York: Bantam: 1980.
- Bateson, Gregory. *Steps to an Ecology of Mind*. San Francisco: Chandler Publishing Co., 1972.
- Bell, Daniel. *The Coming of Post-Industrial Society*. New York: Basic Books, 1973.
- Bender, W. "Imaging and Interactivity." In *Proc. Fifteenth Joint Conference on Image Technology*, Tokyo, November 26, 1984.
- Beniger, James. *The Control Revolution: Technological and Economic Origins of the Information Society*. Cambridge, Massachusetts: Harvard University Press, 1986.
- Bennis, Warren. "Thoughts from a Victim of Info-Overload Anxiety." *Antioch Review*, 1977. Yellow Springs Ohio: Antioch Press, 1977.

- Berlo, David. *The Process of Communication*. New York: Holt, Rinehart and Winston, 1960.
- Blake, Reed and Edwin Haroldser. *A Taxonomy of Concepts in Communication*. New York: Hastings House Publishers, 1975.
- Bolt, Richard. *The Human Interface: Where People and Computers Meet*. Belmont, CA: Lifetime Learning Publications, 1984.
- Bolter, Jay and Michael Joyce. "Hypertext and Creative Writing," In *Hypertext '87 Papers*. Chapel Hill, NC: University of North Carolina, 1987.
- Boorstin, Daniel. *The Republic of Technology: Reflections on Our Future Community*. New York: Harper & Row, 1978.
- Booth-Clibborn, Edward and Daniel Baroni. *The Language of Graphics*. New York: Abrams Publishers, 1979.
- Borenstein, Nathaniel. "Help Texts vs. Help Mechanisms: A New Mandate for Documentation Writers," In *Proceedings of SIGDOC '85*. New York: Association for Computing Machinery, 1985.
- Brand, Stewart. *The Media Lab: Inventing the Future at MIT*. New York: Viking Penguin, 1987.
- Briggs, John. *Fire in the Crucible: The Self-Creation of Creativity and Genius*. Los Angeles: Jeremy P. Tarcher, 1990.
- Brockman, R. John. *Writing Better Computer User Documentation: From Paper to Hypertext*. New York: John Wiley and Sons, 1990.
- Brown, John Seely. *A Personal Perspective on Collaborative Tools*. Austin, Texas: Computer-Supported Cooperative Work Conference, December 1986.

- Bruner, Jerome. *In Search of Mind*. New York: Harper & Row, 1983.
- Bruner, Jerome. *On Knowing: Essays for the Left Hand*. Cambridge, Massachusetts: Harvard University Press, 1962.
- Burnham, David. *The Rise of the Computer State*. New York: Vintage Books, 1984.
- Bush, Vannevar. "As We May Think." *Atlantic Monthly*, July 1945.
- Byers, T. J. "Built by Association." *PC World*, April 1987.
- Campbell, Jeremy. *Grammatical Man: Information, Entropy, Language, and Life*. New York: Simon & Schuster, 1982.
- Communications of the ACM*. Association for Computing Machinery, Vol. 30, No. 1. New York: January 1987.
- Conklin, J. "Hypertext: An Introduction and Survey." *Computer*, September, 1987.
- Dalton, Richard. "Group-Writing Tools: Four That Connect." *Information Week*, March 9, 1987.
- Davis, Stanley M. *Future Perfect*. Reading, Massachusetts: Addison-Wesley, 1987.
- Davis, William and Allison McCormack. *The Information Age*. Reading, Massachusetts: Addison-Wesley, 1979.
- Davison, W. Phillips, James Boylan, and Frederic Yu. *Mass Media: Systems and Effects*, 2nd Edition. New York: Holt, Rinehart & Winston, 1982.
- DeFleur, Melvin and Everette Dennis. *Understanding Mass Communication*, 2nd Edition. Boston: Houghton Mifflin, 1985.
- DeKoven, Bernard. *Power Meetings*, Palo Alto, California: C-EM, 1986.

- DeMarco, Tom and Timothy Lister. *Peopleware: Productive Projects and Teams*. New York: Dorset House, 1987.
- Ditlea, Steve. "Another World: Inside Artificial Reality." *PC Computing*, November 1989.
- Ditlea, Steve. "The Big Link: Artificial Intelligence." *Omni*, September 1987.
- Dondis, Donis. *A Primer of Visual Literacy*. Cambridge, Massachusetts: MIT Press, 1973.
- Donnelly, William. *The Confetti Generation: How the New Communications Technology is Fragmenting America*. New York: Henry Holt, 1986.
- Doyle, Michael and David Strauss. *How to Make Meetings Work*. New York: Jove, 1976.
- Drexler, Eric K. *Engines of Creation: Challenges and Choices of the Last Technological Revolution*. New York: Anchor Books, 1986.
- Dumas, Joseph. *Designing User Interfaces for Software*. Englewood Cliffs, New Jersey: Prentice Hall, 1988.
- Dyson, Esther. "Implementation Support: Beyond the Decision." *Release 1.0.*, August 17, 1988.
- Easterby, R., and H. Zwaga. *The Design and Evaluation of Signs and Printed Material*. New York: Wiley, 1984.
- Engelbart, Douglas C. "A Conceptual Framework for the Augmentation of Man's Intellect," In *Vistas in Information Handling*. edited by Howerton and Weeks. Washington, D.C.: Spartan Books, 1963.
- Engelbart, Douglas C. "A Research Center for Augmenting Human Intellect," In *Proc. FJCC*, vol. 33, no. 1. Montvale, New Jersey: AFIPS Press, Fall 1968.

- Engelbart, Douglas C. "Authorship Provisions in AUGMENT." In *Proc. COMPCON Conference*, San Francisco, February 27-March 1, 1984.
- Engelbart, Douglas C. "Collaborative Support Provisions in AUGMENT." In *Proc. COMPCON Conference*, San Francisco, February 27-March 1, 1984.
- Engelbart, Douglas C. "Coordinated Information Services for a Discipline—or Mission-Oriented Community," In *Proc. Second Annual Computer Communications Conference*, San Jose, California, January 24, 1972.
- Engelbart, Douglas C. "Evolving the Organization of the Future: A Point of View." In *Proc. Stanford International Symposium on Office Automation*, edited by Robert Landau, James Bair, and Jeannie Siegman. Norwood, New Jersey: Ablex Publications, March 23-25, 1980.
- Engelbart, Douglas C. "The Augmented Knowledge Workshop." In *Proc. ACM Conference on the History of Personal Workstations*, Palo Alto, California, January 9-10, 1986.
- Engelbart, Douglas C. "Toward High-Performance Knowledge Workers." In *Proc. AFIPS Office Automation Conference*, San Francisco, April 5-7, 1982.
- Engelbart, Douglas C. "Toward Integrated, Evolutionary Office Automation Systems." In *Proc. Joint Engineering Management Conference*, Denver Colorado, October 16-18, 1978.
- Engelbart, Douglas C. "Workstation History and the Augmented Knowledge Workshop." In *Proc. ACM Conference on the History of Personal Workstations*, Palo Alto, California, January 9-10, 1986.
- Engelbart, Douglas C.; Watson, Richard W.; and Norton, James C. "The Augmented Knowledge Workshop." In *Proc. NCC 73*, 1973.

- Feinstein, Debra. "Colab: Toward a Meeting of Minds." *Benchmark*, Fall 1987.
- Fiderio, Janet. "A Grand Vision: Hypertext Mimics the Brain's Ability to Access Information Quickly and Intuitively by Reference," *Byte*, October 1988.
- Foster, Gregg. "Collaborative Systems and Multi-User Interfaces." University of California, Berkeley, Ph.D. Graduate Thesis, 1986.
- Foster, Gregg and Mark Stefik. "Cognoter, Theory and Practice in a Collaborative Tool." In *Proceedings of the Conference on Computer-Supported Cooperative Work*, New York: ACM, December 1986.
- Fraase, Michael. *Macintosh Hypermedia Volume I: Reference Guide*. Glenview, IL: Scott, Foresman, 1990.
- Fraase, Michael. *Macintosh Hypermedia Volume II: Uses and Implementations*. Glenview, IL: Scott, Foresman, 1990.
- Fraase, Michael. *Macintosh Hypermedia Volume III: Development Tools for HyperCard*. Glenview, IL: Scott, Foresman, 1990.
- Fraase, Michael. *The Ultimate Training Tool: Farallon Computing's MediaTracks*. Homewood, IL: Business One Irwin, 1990.
- Friedhoff, Richard and William Benzon. *Visualization*. New York: Harry Abrams, 1989.
- Frisse, Mark. "From Text to Hypertext," *Byte*, October 1988.
- Gaines, Brian and Mildred Shaw. *The Art of Computer Conversation: A New Medium for Communication*. Englewood Cliffs, New Jersey: Prentice Hall, 1984.
- Garson, Barbara. *The Electronic Sweatshop*. New York: Simon and Schuster, 1988.

- Gaver, W.W. "Auditory Icons: Using Sound in Computer Interfaces," *Human Computer Interaction*, 1986.
- Gerstein, Marc. *The Technology Connection: Strategy and Change in the Information Age*. Reading, Massachusetts: Addison-Wesley, 1987.
- Gibson, William. *Burning Chrome*. New York: Ace Books, 1986.
- Gibson, William. *Count Zero*. New York: Ace Books, 1986.
- Gibson, William. *Mona Lisa Overdrive*. New York: Bantam Books, 1988.
- Gibson William. *Neuromancer*. New York: Ace Books, 1984.
- Gladstein, Deborah. "Groups in Context: A Model of Task Group Effectiveness." *Administrative Science Quarterly*, Volume 29, 1984.
- Goffman, Erving. *Behavior in Public Places: Notes on the Social Organization of Gatherings*. New York: Free Press, 1963.
- Goffman, Erving. *Interaction Ritual*. New York: Doubleday, 1967.
- Goffman, Erving. *The Presentation of Self in Everyday Life*. New York: Doubleday, 1959.
- Greif, Irene, ed. *Computer-Supported Cooperative Work: A Book of Readings*. San Mateo, California: Morgan Kauffman Publishers, 1988.
- Hawken, Paul. *The Next Economy*. New York: Holt, Rinehart & Winston, 1983.
- Hills, P. "Human Communication and Information Technology." In *Trends in Information Transfer*. ed. by P. Hills. Westport, Connecticut: Greenwood Press, 1982.

- Hiltz, Roxanne. *Online Communities: A Case Study of the Office of the Future*. Norwood, New Jersey: Ablex Press, 1984.
- Hirschhorn, Larry. *The Workplace Within: Psychodynamics of Organizational Life*. Cambridge, Massachusetts: MIT Press, 1988.
- Horn, Robert. *Mapping Hypertext*. Lexington, Massachusetts: The Lexington Institute, 1989.
- Horton, William. *Designing and Writing Online Documentation*. New York: John Wiley and Sons, 1990.
- Illich, Ivan. *Tools for Conviviality*. New York: Harper and Row, 1973.
- Innis, Harold. *The Bias of Communication*. Toronto: University of Toronto Press, 1964.
- Innis, Harold. *Empire and Communication*. Toronto: University of Toronto Press, 1972.
- Irby, Charles H. "Display Techniques for Interactive Text Manipulation." In *Proc. NCC 74*, 1974.
- Janis, Irving. *Victims of Groupthink*, 2nd ed. Boston: Houghton Mifflin, 1982.
- Johansen, Robert. *Groupware: Computer Support for Business Teams*. New York: Free Press, 1988.
- Johansen, Robert, Jacques Vallee, and Kathleen Vian. *Electronic Meetings*. Reading, Massachusetts: Addison-Wesley, 1979.
- Johansen, Robert. *Teleconferencing and Beyond*. New York: McGraw-Hill, 1984.
- Johnson-Laird, P.H. *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Cambridge, Massachusetts: Harvard University Press, 1983.

- Johnson-Lenz, P. and T. Johnson-Lenz. "Groupware: The Emerging Art of Orchestrating Collective Intelligence." Presented at the First Global Conference on the Future, Toronto, 1980.
- John-Steiner, Vera. *Notebooks of the Mind*. New York: Harper & Row, 1985.
- Kay, Alan. "Computer Software," *Scientific American*. March 1984.
- Kay, Alan & Goldberg, Adele. "Personal Dynamic Media." *Computer*, March 1977.
- Keen, Peter and Michael S. Scott Morton. *Decision Support Systems*. Reading Massachusetts: Addison-Wesley, 1978.
- Kelly, John. "Technology That Supports Meetings." *Patricia Seybold's Office Computing Report*, September 1988.
- Kemp, C. Gratton. *Perspectives on the Group Process*. Boston: Houghton Mifflin, 1964.
- Kerr, Elaine, and Starr Roxanne Hiltz. *Computer-Mediated Communication Systems*. New York: Academic Press, 1982.
- Klapp, Orrin. *Overload and Boredom: Essays on the Quality of Life in the Information Society*. New York: Greenwood Press, 1986.
- Kobayashi, Koji. *Computers and Communications*. Cambridge, Massachusetts: MIT Press, 1986.
- Koestler, Arthur. *The Act of Creation*. New York: Macmillan, 1964.
- Krueger, Myron. *Artificial Reality*. Reading, Massachusetts: Addison-Wesley, 1985.
- Lancaster, F.W. *Toward Paperless Information Systems*. New York: Academic Press, 1978.

- Laurel, Brenda, ed. *The Art of Human-Computer Interface Design*. Reading, Massachusetts: Addison-Wesley, 1990.
- Lederman, Linda Costigan. "Communication in the Workplace: The Impact of the Information Age and High Technology on Interpersonal Communication in Organizations." *Intermedia: Interpersonal Communication in a Media World*, 3rd Edition. Gary Gumpert and Robert Cathcard, eds. New York: The Oxford University Press, 1986.
- Levine, Howard and Howard Rheingold. *The Cognitive Connection: Thought and Language in Man and Machine*. New York: Prentice Hall Press, 1987.
- Licklider, J.C.R., Robert Taylor, and E. Herbert. "The Computer as a Communication Device." In *International Science and Technology*, April 1978.
- Linstone, Harold and Murray Turoff. *The Delphi Method: Techniques and Applications*. Reading, Massachusetts: Addison-Wesley, 1975.
- Luft, Joseph. *Group Processes*. Palo Alto, California: National Press Books, 1963.
- McLuhan, Marshall and Eric McLuhan. *Laws of Media: The New Science*. Toronto: University of Toronto Press, 1988.
- McLuhan, Marshall and Bruce Powers. *The Global Village: Transformations in World Life and Media in the 21st Century*. New York: Oxford University Press, 1989.
- McLuhan, Marshall. *The Gutenberg Galaxy: The Making of Typographic Man*. Toronto: University of Toronto Press, 1962.
- McLuhan, Marshall. *Understanding Media: The Extensions of Man*. New York: McGraw-Hill, 1965.

- Meadow, C.T., and P.A. Cochrane. *Basics of Online Searching*. New York: John Wiley and Sons, 1981.
- Meyrowitz, Joshua. *No Sense of Place: The Impact of Electronic Media on Social Behavior*. New York: Oxford University Press, 1985.
- Meyrowitz, Norman. & Van Dam, A. "Interactive Editing Systems." *Computing Surveys*, March 1982.
- Miller, George. *The Psychology of Communication*. New York: Basic Books, 1967.
- Minsky, Marvin. "A Framework for Representing Knowledge." In *The Psychology of Computer Vision*, ed. by P. Winston, New York: McGraw-Hill, 1975.
- Minsky, Marvin. *The Society of Mind*. New York: Simon & Schuster, 1986.
- Mumford, Louis. *Technics and Civilization*. New York: Harcourt Brace Jovanovich, 1934.
- Naisbitt, John. *Megatrends*. New York: Warner Books, 1982.
- Negroponte, Nicholas. *The Architecture Machine*. Cambridge, Massachusetts: MIT Press, 1970.
- Negroponte, Nicholas. *Soft Architecture Machines*. Cambridge, Massachusetts: MIT Press, 1975.
- Nelson, Ted. "A Conceptual Framework for Man-Machine Everything." In *Proc. National Joint Computer Conference*, 1973.
- Nelson, Ted. "A Dream for Irving Snerd." *Creative Computing*, July 1977.
- Nelson, Ted. "A File Structure for the Complex, the Changing and the Indeterminate." In *Proc. Association for Computing Machinery*, 1965.

- Nelson, Ted. "A New Home for the Mind." *Datamation*, March 1982.
- Nelson, Ted. "A Vision of the Future." *Publishers Weekly*, 23 November 1986.
- Nelson, Ted. "As We Will Think." In *Proc. Online 72 Conference*, Brunel U., Uxbridge England.
- Nelson, Ted. "Barnum-Tronics." *Swarthmore College Alumni Bulletin*, December 1970.
- Nelson, Ted. *Computer Lib/Dream Machines*, rev. ed. Redmond, Washington: Tempus Books of Microsoft Press, 1987.
- Nelson, Ted. "Computopia and Cybercrud." in *Computers in Instruction*, edited by Levien. The Rand Corporation, 1974.
- Nelson, Ted. "Computopia Now!" In *Digital Deli*, edited by Steve Ditlea. San Francisco: Workman Publishing, 1984.
- Nelson, Ted. "Data Realms and Magic Windows." In *Proc. ACPA-5*, Meeting of the Assn. of Computer Programmers and Analysts, 1975.
- Nelson, Ted. "Electronic Publishing and Electronic Literature." In *Information Technology in Health Science Education*, edited by DeLand. Plenum Press, 1978.
- Nelson, Ted. "Getting It Out of Our System." In *Critique of Information Retrieval*, edited by Schechter. Thompson Books, 1967.
- Nelson, Ted. "Hypertext Notes." Ten brief essays on hypertext forms, ca. 1966. Manuscript.
- Nelson, Ted. "Interactive Systems and the Design of Virtuality." *Creative Computing*, November 1980.

- Nelson, Ted. "Las Vegas Confrontation Sit-Out: A CAI Radical's View from Solitary." *SIGCUE Newsletter*, 1971.
- Nelson, Ted. *Literary Machines*, electronic ed. Bellevue, Washington: OWL International Inc., 1987.
- Nelson, Ted. "New Media and Creativity Systems." Graphical brochure intended to expound computer graphics and related concepts, ca. 1966.
- Nelson, Ted. "No More Teacher's Dirty Looks." *Computer Decisions*, September 1970.
- Nelson, Ted. "On the Road to Xanadu." *Personal Computing*, November 1987.
- Nelson, Ted. "Replacing the Printed Word." In *Proc. IFIP 80*, edited by S.H. Lavington. North-Holland Publishing Co., 1980.
- Nelson, Ted. "Stretchtext." Unpublished paper on stretchable hypertext form, 1966. Photocopy.
- Nelson, Ted. "Suggestion for an On-Line Braille Display." In *Proc. Society for Information Display*, Fall 1965.
- Nelson, Ted. "The Hypertext." In *Proc. World Documentation Federation*, 1965.
- Nelson, Ted. "The Magicians, the Snark and the Camel." *Creative Computing*, December 1981.
- Nelson, Ted. "The Tyranny of the File." *Datamation*, 15 December 1986.
- Nelson, Ted; Carmody, Steven; and et. al. "A Hypertext Editing System for the 360." In *Pertinent Concepts in Computer Graphics*, edited by Faiman and Nievergelt. U. Illinois Press, 1969.

- Nelson, Ted; DeFanti, Tom; and Sandin, Dave. "Computer Graphics as a Way of Life." In *Proc. Siggraph Conference*, 1974.
- Norman, Donald A. and Stephen W. Draper (eds.). *User Centered Design: New Perspectives on Human-Computer Interaction*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1986.
- Norman, Donald A. *The Psychology of Everyday Things*. New York: Basic Books, 1988.
- Olson, Margrethe. "Computer-Supported Cooperative Work." *Office: Technology and People*, Volume 3, Number 2, August 1987.
- Ong, W. *Orality and Literacy: The Technologizing of the Word*. New York: Methuen, 1982.
- Onosko, Timothy. "Let There Be Light," *Creative Computing*, September 1985.
- Oren, Tim. "The Architecture of Static Hypertexts," In *Hypertext '87 Papers*. Chapel Hill, NC: University of North Carolina, 1987.
- Patton, Bobby and Kim Griffin. *Problem Solving Group Interaction*. New York: Harper & Row, 1973.
- Pool, Ithiel de Sola. *Technologies of Freedom: On Free Speech in an Electronic Age*. Cambridge, Massachusetts: Harvard University Press, 1983.
- Proceedings of the 1986 Conference on Computer-Supported Work*. Austin, Texas: Conference Committee for CSCW, 1986.
- Proceedings of the Conference on Computer-Supported Cooperative Work, September 1988*. Portland, Oregon: Association for Computing Machinery, 1988.

- Proceedings of the Symposium on Technological Support for Work Group Collaboration*. Center for Research on Information Systems, Graduate School of Business, New York University, 1987.
- Rheingold, Howard. *Tools for Thought: The People and Ideas behind the Next Computer Revolution*. New York: Simon and Schuster, 1985.
- Rockley, Ann and Gord Graham. "Hypermedia—A Web of Thought," In *Proceedings of 35th International Technical Communication Conference*. Washington DC: Society for Technical Communication, 1988.
- Rogers, Everett and Rekha Agarwala-Rogers. *Communication in Organizations*. New York: Free Press, 1976.
- Rosenberg, Richard. *Computers and the Information Society*. New York: Wiley, 1986.
- Roszak, Theodore. *The Cult of Information*. New York: Pantheon Books, 1986.
- Schrage, Michael. *Shared Minds: The New Technologies of Collaboration*. New York: Random House, 1990.
- Shannon, Claude. "A Mathematical Theory of Information." In *Bell System Technical Journal*, Volume 27, 1948.
- Shannon, Claude and Warren Weaver. *The Mathematical Theory of Communication*. Champaign-Urbana, Illinois: The University of Illinois Press, 1949.
- Shneiderman, Ben. *Designing the User Interface*. Reading, Massachusetts: Addison-Wesley Publishing Company, 1987.
- Smith, John and Stephen Weiss. "Hypertext," *Communications of the ACM*. July 1988.

- Smith, John, Stephen Weiss, and Gordon Ferguson. "A Hypertext Writing Environment and Its Cognitive Basis," In *Hypertext '87 Papers*. Chapel Hill, NC: University of North Carolina, 1987.
- Stanley, Robert. *Mediavisions: The Art and Industry of Mass Communication*. New York: Praeger, 1987.
- Stefik, Mark et al. "Beyond the Chalkboard," *Communications of the ACM*, January 1987.
- Stefik, Mark. "The Next Knowledge Medium." *AI Magazine*. Volume 7, Number 1, Spring 1986.
- Stefik, Mark and John Seely Brown. "Toward Portable Ideas," *Communications of Xerox PARC*. Xerox Corporation, Palo Alto Research Center, 1988.
- Stefik, Mark et al. "WYSIWIS Revised: Early Experiences with Multiuser Interfaces." *ACM Transactions on Office Information Systems*, Volume 5, Number 2, April 1987.
- Steiner, Ivan. *Group Process and Productivity*. New York: Academic Press, 1972.
- Stillman, Stanley, ed. *The Information Industry: Structure, Characteristics, and Trends*. New York: Access Publishing, 1987.
- Strassman, Paul. *Information Payoff: The Transformation of Work in the Electronic Age*. New York: The Free Press, 1985.
- Technological Support for Work Group Collaboration*. 1987 NYU Symposium, Center for Research on Information Systems, New York University, 1987.
- Toffler, Alvin. *Powershift: Knowledge, Wealth, and Violence at the Edge of the 21st Century*. New York: Bantam, 1990

- Toffler, Alvin. *The Third Wave*. New York, New York: Morrow, 1980.
- Trigg, R.; Suchman, L.; and Halasz, F. "Supporting Collaboration in Notecards." In *Proc. Conference on Computer-Supported Cooperative Work '86*, Austin, Texas, December 1986.
- Tufte, Edward. *The Visual Display of Quantitative Information*. Cheshire, Connecticut: Graphics Press, 1984.
- Tufte, Edward. *Envisioning Information*. Cheshire, Connecticut: Graphics Press, 1990.
- Turoff, Murray. "Delphi Conferencing: Computer-Based Conferencing with Anonymity." In *Technological Forecasting and Social Change*, 1972.
- Turoff, Murray and Roxanne Hiltz. *The Network Nation: Human Communication via Computer*. Reading, Massachusetts: Addison-Wesley, 1978.
- Vallee, Jacques. *Computer Message Systems*. New York: McGraw-Hill, 1984.
- Vallee, Jacques. *The Network Revolution: Confessions of a Computer Scientist*. Berkeley, California: And/Or Press, 1982.
- Van Dam, A.; Feiner, S.; and Nagy, S. "An Experimental System for Creating and Presenting Interactive Graphical Documents." *Trans. Graphics*, vol. 1, no. 1, 1982.
- Weyer, S.A. *Searching for Information in a Dynamic Book*. SCG-82-1, Palo Alto, California: Xerox Palo Alto Research Center, February 1982.
- Whorf, Benjamin. *Language, Thought, and Reality*. Cambridge, Massachusetts: MIT Press, 1956.

- Widmer, Kingsley. "Sensibility Under Technocracy." *Human Connection and the New Media*. Barry Schwartz, ed. Englewood Cliffs, New Jersey: Prentice-Hall, 1973.
- Winograd, Terry and Fernando Flores. *Understanding Computers and Cognition*. Norwood, New Jersey: Ablex, 1986.
- Wurman, Richard Saul. "Hats." *Design Quarterly* 145. Mildred Friedman, ed. Cambridge, Massachusetts: The MIT Press, 1989.
- Wurman, Richard Saul. *Information Anxiety*. New York: Doubleday, 1989.
- Yankelovich, N.; Haan, B.; Meyrowitz, N.; and Drucker, S. "Intermedia: The Concept and the Construction of a Seamless Information Environment." *IEEE Computer*, January 1988.
- Zuboff, Shoshana. *In the Age of the Smart Machine*. New York: Basic Books, 1988.

Appendix B

Product Sources

AppleShare
Apple Computer Inc.
20525 Mariani Avenue
Cupertino, CA 95014
408/996-1010

AppleShare Print Server
Apple Computer Inc.
20525 Mariani Avenue
Cupertino, CA 95014
408/996-1010

CheckNet
Farallon Computing Inc.
2000 Powell Street
Suite 600
Emeryville, CA 94608
415/596-9100

DataClub
International Business Software Inc.
1230 Oakmead Parkway
Suite 310
Sunnyvale, CA 94088
408/720-0170

GatorBox
Cayman Systems Inc.
26 Landsdowne Street
Cambridge, MA 02139
617/494-9270

GatorShare
Cayman Systems Inc.
26 Landsdowne Street
Cambridge, MA 02139
617/494-9270

GoFER
Microlytics, Inc.
300 Main Street
Suite 801
East Rochester, NY 14445
716/377-0130

GraceLAN
Technology Works Inc.
4030 Braker Lane West
Suite 350
Austin, TX 78759
512/794-8533

Inspiration 2.0
Ceres Software Inc.
2520 SW Hamilton Street
Portland, OR 97201
503/245-9011

MacLink Plus
DataViz Inc.
35 Corporate Drive
Trumbull, CT 06611
203/268-0030

MarkUp

Mainstay Software
5311-B Derry Avenue
Agoura Hills, CA 91301
818/991-6540

Microsoft Mail

Microsoft Corporation
16011 NE 36th Way
P.O. Box 97017
Redmond, WA 98073
206/882-8080

NetAtlas

Farallon Computing Inc.
2000 Powell Street
Suite 600
Emeryville, CA 94608
415/596-9100

NetUpdater

MDG Computer Services
634 South Dunton Avenue
Arlington Heights, IL 60005
708/818-9991

On Location

ON Technology Inc.
155 Second Street
Cambridge, MA 02141
617/876-0900

Personal Server Network

Information Presentation Technologies Inc.
5000 North Parkway Calabasas Road
Suite 304
Calabasas, CA 91302
800/233-9993
818/347-7791

PhoneNET
Farallon Computing Inc.
2000 Powell Street
Suite 600
Emeryville, CA 94608
415/596-9100

PhoneNET Liaison
Farallon Computing Inc.
2000 Powell Street
Suite 600
Emeryville, CA 94608
415/596-9100

PhoneNET Card PC•LocalTalk
Farallon Computing Inc.
2000 Powell Street
Suite 600
Emeryville, CA 94608
415/596-9100

QuickMail
CE Software
1854 Fuller Road
P.O. Box 65580
West Des Moines, IA 50265

Retrospect Remote
Dantz Development Corporation
1510 East Walnut Street
Berkeley, CA 94709
415/849-0293

TeleFinder
Spider Island Software
4790 Irvine Blvd.
Suite 105-347
Irvine, CA 92720
714/669-9260

Timbuktu
Farallon Computing Inc.
2000 Powell Street
Suite 600
Emeryville, CA 94608
415/596-9100

Timbuktu/Remote
Farallon Computing Inc.
2000 Powell Street
Suite 600
Emeryville, CA 94608
415/596-9100

Timeslips Network Edition
Timeslips Corporation
239 Western Avenue
Essex, MA 01929
508/768-6100

Traffic Watch II
Farallon Computing Inc.
2000 Powell Street
Suite 600
Emeryville, CA 94608
415/596-9100

Appendix C

Glossary

accelerator card: An expansion card that contains another processor that shares the work normally performed only by the computer's main microprocessor. An accelerator card speeds up processing time.

access method: The rules that manage how all the computers and other devices on a network can send information through the same physical medium in an orderly fashion.

access privileges: The privileges to open and make changes to folders and their contents; they are given to or withheld from users. By setting access privileges, you can control access to confidential information stored in folders on a server.

Access Privileges window: When using the AppleShare file server, a window that displays the access privileges, owner, group, and other information about a folder or volume. You use the Get Privileges command in the File menu of the Finder or the Access Privileges desk accessory in the Apple menu to display the window and to review, set, or change access privileges for a folder or volume.

ACIA: Abbreviation for Asynchronous Communications Interface Adapter; a type of communications IC used in some

Apple computers. An ACIA converts data from parallel to serial form and vice versa. It handles serial transmission and reception and RS-232-C signals under the control of its internal registers, which can be set and changed by firmware or software. Compare SCC.

acronym: A word formed from the initial letter or letters of the main parts of a compound term, such as ROM (from read-only memory or Fortran (from Formula Translator).

activate: To make a nonactive window active by clicking anywhere inside it.

active star topology: A star topology with the controller device located at the center of the star. The controller interconnects multiple branch cables and actively repeats network signals.

active window: The frontmost window on the desktop; the window where the next action will take place. An active window's title bar is highlighted.

ADB: See **Apple Desktop Bus**.

addressing: A scheme, determined by network protocols, for identifying the sending device and destination device for any given item of information traveling on a network.

administrator: The person who sets up a file server, registers users and their passwords, creates AppleShare groups, and maintains the server.

ALAP: See **AppleTalk Link Access Protocol**.

alert: A warning or report of an error in the form of an alert box, a sound from the computer's speaker, or both. See also **alert box**.

alert box: A box that appears on the screen to give a warning or to report an error message. Its appearance is usually accompanied by a sound warning such as a beep.

alias: (n.) An alternate name used to invoke or identify a command, a network host, a list of users, or some other applicable entity. (v.) To provide an entity with an alternate name.

American Standard Code for Information Interchange:
See ASCII.

analog: (adj.) Varying smoothly and continuously over a range, rather than changing in discrete jumps. For example, a conventional 12-hour clock face is an analog device that shows the time of day by the continuously changing position of the clock's hands. Compare **digital**.

analog data: Data in the form of continuously variable quantities. Compare **digital data**.

analog signal: A signal that varies continuously over time, rather than being sent and received in discrete intervals. Compare **digital signal**.

analog-to-digital converter (ADC): A device that converts quantities from analog to digital form. For example, computer hand controls convert the position of the control dial (an analog quantity) into a discrete number (a digital quantity) that changes in steps even when the dial is turned smoothly.

analog transmission: Transmission of a continuously variable signal as opposed to a discretely variable signal. Compare **digital transmission**.

ANSI: Acronym for American National Standards Institute, which sets standards for many technical fields and provides the most common standard for computer terminals.

Apple Desktop Bus (ADB): A low-speed, input-only serial bus with connectors on the back panel of the computer that you use to attach the keyboard, mouse, and other Apple Desktop Bus devices, such as graphics tablets, hand controls, and specialized keyboards.

Apple menu: The menu farthest to the left in the menu bar, indicated by an Apple symbol, from which you choose desk accessories.

AppleTalk Link Access Protocol (ALAP): The lowest-level protocol in the AppleTalk architecture, managing node-to-node delivery of frames on a single AppleTalk network.

AppleTalk Manager: An interface to a set of device drivers that enable programs to send and receive information via an AppleTalk network.

AppleTalk network system: The system of network software and hardware used in various implementations of Apple's communications network.

AppleTalk protocol: Apple Computer Inc.'s network standard that defines how devices communicate with each other across a network.

AppleTalk Transaction Protocol (ATP): An AppleTalk protocol that's a Datagram Delivery Protocol (DDP) client. It allows one ATP client to request another ATP client to perform some activity and to report the activity's result as a response to the requesting socket with guaranteed delivery. See also **Datagram Delivery Protocol**.

application (usually application program): Software that a user interacts with, or uses. Examples of applications are word processing programs, spreadsheet programs, and electronic mail programs.

architecture: The design plan of a network that determines how the network's components function together.

archive: (n.) (1) A collection of object files, plus a table of contents. (2) Any collection of files saved simultaneously for backup purposes, usually intended for longer storage than are daily backups. Compare **backup**. (v.) To save a collection of files for storage. Compare **back up**.

ARPANET: A wide area network that links government, academic, and industrial installations around the world. Primarily connecting research sites, the ARPANET was developed in the 1960s by the Advanced Research Projects Agency of the U.S. Department of Defense. See also **Defense Data Network**.

ASCII: Acronym for American Standard Code for Information Interchange (pronounced "ASK-ee"). A standard that assigns a unique binary number to each text character and control character. ASCII code is used for representing text inside a computer and for transmitting text between computers or between a computer and a peripheral device. Compare EBCDIC. See also **high ASCII characters, low ASCII characters**.

asynchronous: Not synchronized by a mutual timing signal or clock. Compare **synchronous**.

asynchronous communication: See **asynchronous transmission**.

Asynchronous Communications Interface Adapter: See **ACIA**.

asynchronous execution: A method of routine execution that leaves a program free to perform other tasks until the routine is completed.

asynchronous I/O: The capability to perform an I/O operation while its calling process continues to run. With synchronous I/O, the calling process "sleeps" until the I/O operation is finished.

asynchronous transmission: A method of data transmission in which the receiving and sending devices don't share a common timer, and no timing data is transmitted. Each information character is individually synchronized, usually by the use of start and stop bits. The time interval between

characters isn't necessarily fixed. Compare **synchronous transmission**.

ATP: See **AppleTalk Transaction Protocol**.

backbone topology: A network topology with devices connected to a single, continuous cable. Also called a **trunk topology**.

background activity: A program or process that runs while the user is engaged with another application.

background printing: Printing from one application while another application is running.

background processing: In multitasking environments, the operating system's ability to process lower-priority tasks while you perform other work on the computer.

backup: (n.) A copy of a disk or of a file on a disk. It's a good idea to make backups of all your important disks and to use the copies for everyday work, keeping the originals in a safe place. (Some program or startup disks cannot be copied.) Compare **archive**.

back up: (v.) To make a spare copy of a disk or of a file on a disk. Backing up your files and disks ensures that you won't lose information if the original is lost or damaged. Compare **archive**.

backup bit: A bit in a file's access byte that tells backup programs whether the file has been altered since the last time it was backed up.

bandwidth: The range of transmission frequencies that a network can use. The greater the bandwidth, the greater the amount of information that can travel on the network at one time.

baseband: A transmission method in which a network uses its entire transmission frequency range to send a single communication or signal. Compare **broadband**.

baud: (1) A unit of data transmission speed: the number of discrete signal-state changes (signal events) per second. Often, but not always, equivalent to **bits per second**. Compare **bit rate**. (2) The maximum speed at which data can be sent down a channel, such as a telephone line; often confused with the actual speed at which the data is transmitted between two computers, measured in bits per second.

BBS: See **bulletin board system**.

binary synchronous communication (BSC): A type of protocol developed by IBM that uses synchronization of characters to control the transfer of data over communication lines. Also referred to as **bi-sync communication**. Compare **SNA/SDLC**.

bit: A contraction of **binary digit**. The smallest unit of information that a computer can hold. The value of a bit (1 or 0) represents a simple two-way choice, such as yes or no, on or off, positive or negative, something or nothing.

bit rate: The speed at which bits are transmitted, usually expressed as **bits per second**, or **bps**. Compare **baud**.

bits per second: See **bit rate**.

boot block: (1) An area on a formatted disk that signals the computer that the disk contains an application to be started up. (2) The first block of a file system, or the first two logical blocks of a volume. The boot block contains the system's startup instructions.

boot device: The peripheral device that reads an operating system's initial startup instructions.

boot disk: See **startup disk**.

bps (bits per second): See **bit rate**.

bridge: (1) A device that lets you connect two or more networking systems together. (2) A combination of hardware and software that connects two or more networks in an internet. Bridges are used to increase the number of devices and the distances covered in a network. See also **internet**, **zone**.

broadband: A transmission method in which the network's range of transmission frequencies is divided into separate channels and each channel is used to send a different signal. Broadband transmission is often used to send signals of different kinds simultaneously, such as voice and data. Contrast with **baseband**.

bulletin board system (BBS): A computerized version of the bulletin boards frequently found in grocery stores—places to leave messages and to advertise things you want to buy or sell. One thing you get from a computerized bulletin board that you can't get from a cork board is free software.

bus topology: A layout scheme in which devices on a network are connected along the length of a main cable, or **bus**, rather than in a daisy chain or a loop.

byte: The number of bits used to represent a character. For personal computers, a byte is usually eight bits.

cable: An insulated bundle of wires with connectors on the ends. Examples are serial cables, disk drive cables, and LocalTalk cables.

cable extender: A small plastic adapter with a LocalTalk socket on either end that allows you to connect two LocalTalk cables together.

card: (1) A printed-circuit board that plugs into one of the computer's expansion slots. (2) A printed-circuit board or card connected to the bus in parallel with other cards. Also called a peripheral card, a device, or a module.

carrier: The background signal on a communication channel that is modified to carry information. Under RS-232-C rules, the carrier signal is equivalent to a continuous MARK (1) signal; a transition to 0 then represents a start bit.

carrier sensing: An access method, often referred to by the acronym CSMA, for carrier sense, multiple access. According to this method, computers check the network medium to see if it is being used and wait for it to be free before transmitting.

CCITT: Abbreviation for Consultative Committee on International Telegraphy and Telephony; an international committee that sets standards and makes recommendations for international communication.

choose: To pick a command by dragging through a menu. You often choose a command after you've selected something for the program to act on; for example, selecting a disk and choosing the Open command from the File menu.

Chooser: A desk accessory that lets you configure your computer system to print on any printer for which there's a printing resource on the current startup disk. If you're part of an AppleTalk network system, you use the Chooser to connect and disconnect from the network and choose among devices connected to the network. You can also specify a user name that the system uses from time to time—when you're printing on a LaserWriter, for example.

Chooser-level driver: The software that allows you to select "Print" from an application's menu to produce output. The alternative is a separate, more cumbersome software output utility.

circuit switching: A transmission method in which the path, or circuit, along the network between the sending and receiving devices is dedicated to transmitting data only between these two devices. Compare with packet switching.

Clear To Send: An RS-232-C signal from a DCE to a DTE that is normally kept false until the DCE makes it true, indicating that all circuits are ready to transfer data. See also **Data Communication Equipment, Data Terminal Equipment.**

click: (v.) To position the pointer on something, and then press and quickly release the mouse button. (n.) The act of clicking.

client: A computer that has access to services on a network. The computers that provide services are called **servers**. A user at a client may request file access, remote log-in, file transfer, printing, or other available services from servers.

Clipboard: The holding place for what you last cut or copied; a buffer area in memory. Information on the Clipboard can be inserted (pasted) into documents.

closed driver: A device driver that cannot be read from or written to.

coaxial cable: One type of media used in local area networks. Consists of a central wire surrounded by a layer of insulation, a conductive metal shield, and an outer insulator.

Command key: A key that, when held down while another key is pressed, causes a command to take effect. When held down in combination with dragging the mouse, the Command key lets you drag a window to a new location without activating it. The Command key is marked with a propeller-shaped symbol. On some machines, the Command key has both the propeller symbol and the Apple symbol on it.

communication controller: A type of communication control unit with its operations controlled by a program stored and executed at that unit. For example, the IBM 3704 and 3705 are communication controllers.

communication mode: An operating state in which a serial card or port is prepared to exchange data and signals with a

modem or other type of data communication equipment (DCE).

communications protocol: See **protocol**.

configuration: (1) A general-purpose computer term that can refer to the way you have your computer set up. (2) The total combination of hardware components—central processing unit, video display device, keyboard, and peripheral devices—that make up a computer system. (3) The software settings that allow various hardware components of a computer system to communicate with one another.

configure: To change software or hardware actions by changing settings. Configurations can be set or reset in software or by manipulating hardware jumpers, switches, or other elements.

connect time: The amount of time you spend connected to an information service.

context sensitive: Able to perceive the situation in which an event occurs. For example, an application program might present help information specific to the particular task you're performing, rather than a general list of commands; such help would be context sensitive.

Control Panel: A desk accessory that lets you change the speaker volume, the keyboard repeat speed and delay, mouse tracking, color display, and other features.

coprocessor: An auxiliary processor that is designed to relieve the demand on the main processor by performing a few specific tasks. Coprocessors may favor a certain set of operations, like floating-point calculations for graphics instruction looping, and therefore they can optimize the speed at which such operations are processed. Generally, coprocessors handle tasks that could be performed by the main processor running appropriate software but would be performed much more slowly that way.

cut and paste: To move something from one place in a document to another in the same document or a different one. It's the computer equivalent of using scissors to clip something and glue to paste the clipping somewhere else.

daisy chain: (n.) A colloquial term for a group of devices connected to a host device, where the first device in the "chain" is connected to the host, the second device is connected to the first, the third device is connected to the second, and so on. (v.) To link together sequentially.

daisy chain topology: A network topology where each device on the network is directly connected to the next with a modular extension cable and a connecting device.

data: Information, especially information used or operated on by a program. The smallest unit of information a computer can understand is a bit.

data bits: In the stream of bits being sent from your computer to a peripheral device or another computer, the bits that contain meaningful information; distinguished from bits used to indicate that a character is about to start, has stopped, or is correct.

data bus: The path along which general information is transmitted within the computer. The wider the data bus, the more information can be transmitted at once. The Macintosh II, for example, has a 32-bit data bus. Thus, 32 bits of information can be transferred at a time, so that information is transferred twice as fast as in 16-bit computers (assuming equal system clock rates).

Data Carrier Detect (DCD): An RS-232-C signal from a DCE (such as a modem) to a DTE (such as a Macintosh) indicating that a communication connection has been established. See also **Data Communication Equipment, Data Terminal Equipment.**

Data Communication Equipment (DCE): As defined by the RS-232-C standard, any device that transmits or receives information. Usually this device is a modem.

Datagram Delivery Protocol (DDP): An AppleTalk protocol managing socket-to-socket delivery of datagrams over a networking system.

Data Set Ready (DSR): An RS-232-C signal from a DCE to a DTE indicating that the DCE has established a connection. See also **Data Communication Equipment, Data Terminal Equipment.**

Data Terminal Equipment (DTE): As defined by the RS-232-C standard, any device that generates or absorbs information, thus acting as an endpoint of a communication connection. A computer might serve as a DTE.

Data Terminal Ready (DTR): (1) One of the handshake lines in a data transmission interface. See also **hardware handshake.** (2) An RS-232-C signal from a DTE to a DCE indicating a readiness to transmit or receive data. See also **Data Communication Equipment, Data Terminal Equipment.**

DCD: See **Data Carrier Detect.**

DCE: See **Data Communication Equipment.**

DDN: See **Defense Data Network.**

DDP: See **Datagram Delivery Protocol.**

default: A value, action, or setting that a computer system assumes, unless the user gives an explicit instruction to the contrary. Default values prevent a program from stalling or crashing if no value is supplied by the user.

Defense Data Network (DDN): A single, wide area, packet-switching network that integrates the ARPANET research network and the MILNET defense network.

desktop: Your working environment on the computer—the menu bar and the gray area on the screen. You can have a number of documents on the desktop at the same time. At the Finder level, the desktop displays the Trash and the icons (and windows) of disks that have been accessed.

desktop environment: A set of program features that make user interactions with an application resemble the way people work on a desk top. Commands appear as options in pull-down menus, and material being worked on appears in areas of the screen called windows. The user selects commands or other material by using the mouse to move a pointer around on the screen or by using keyboard equivalents.

device: (1) A hardware component of a computer system, such as a video monitor, a disk drive, or a printer. Also called a **peripheral device** because such equipment is often physically separate from, but attached to, the computer. (2) A part of the computer, or a piece of external equipment, that can transfer information. (3) Any piece of equipment that can be attached to a network—a computer, a printer, a file server, a print server, or any other peripheral device.

dialog box: (1) A box that contains a message requesting more information from you. Sometimes the message warns you that you're asking your computer to do something it can't do or that you're about to destroy some of your information. In these cases, the message is often accompanied by a beep. (2) A box that a Macintosh application displays to request information or to report that it is waiting for a process to complete.

digital: Represented in a discrete (noncontinuous) form, such as numerical digits or integers. For example, contemporary digital clocks show the time as a digital display (such as 2:57) instead of using the positions of a pair of hands on a clock face. Compare **analog**.

digital data: Data that can be represented by digits—that is, data that are discrete rather than continuously variable. Compare **analog data**.

digital transmission: Transmission of a discretely variable signal as opposed to a continuously variable signal. Quantities such as temperature are continuously variable—that is, a given temperature value may fall between 2 degrees—and so are described as analog. Data characters, on the other hand, are coded in discrete, separate pulses or signal levels—either 0 or 1, on or off, and so on—and are referred to as digital. Compare **analog transmission**.

document: Whatever you create with an application program—information you enter, modify, view, or save. See also **file**.

document window: The window that displays a document image or a document opened from disk.

download: To transfer files or information from one computer to another, or from a computer to a peripheral device such as a printer. A printer will download fonts if a user prints a document containing fonts that are stored on a Macintosh computer but not stored in the printer's memory.

downloadable font: A font that can be downloaded into a printer or computerized typesetter, which means that tables telling how to construct the type characters are sent from the computer to the output device. By accepting additional character sets—downloadable fonts—an output device can produce a greater number and variety of typefaces. To be able to accept downloadable fonts, a printer or typesetter must have sufficient memory and processing power to receive and store the images.

driver: (1) A program, usually in a System Folder, that lets a peripheral device and a computer send and receive files. Printer drivers control printers; a hard disk driver controls exchanges between a hard disk and a computer.

DSR: See Data Set Ready.

DTE: See Data Terminal Equipment.

DTR: See Data Terminal Ready.

dynamic addressing: AppleTalk's method of assigning network node addresses to each device as it is started up. Any device can have a different network node address each time it is started up.

EBCDIC: Acronym for Extended Binary-Coded Decimal Interchange Code (pronounced "EB-si-dik"). A code used by IBM that represents each letter, number, special character, and control character as an 8-bit binary number. EBCDIC has a character set of 256 8-bit characters. Compare ASCII.

echo: To send an input character back to the originating device for display or verification; for example, to send each character of your message back to your monitor so you know it's been sent to another computer.

electronic mail: A network service that enables users to send and receive messages via computer.

ergonomics: The science of designing work environments that allow people and things to interact efficiently and safely. Sometimes called **human engineering**.

Ethernet: A high-speed local area network that consists of a cable technology and a series of communication protocols. The hardware (cable) provides the physical link to connect systems together. The TCP/IP protocol allows different computers to exchange information over a network. The Ethernet specification was developed by Digital Equipment Corporation, Intel Corporation, and Xerox Corporation. Ethernet is a registered trademark of Xerox Corporation.

Ethernet backbone: A network topology with separate AppleTalk networks interconnected to Ethernet gateways that

are, in turn, connected to a single, continuous Ethernet cable.

Ethernet cable system: A system of high-performance coaxial cables widely used in the communications industry. Ethernet cables can be part of an AppleTalk network system.

EtherTalk: A high-speed AppleTalk network system that uses the cables of an Ethernet network. Ethernet is a widely used communications network. The Apple software and interface card that allow AppleTalk networking protocols to run on high-speed Ethernet media at 10 megabits per second—more than 40 times faster than the standard rate supported by Macintosh computers of 230.4 kilobits per second.

even parity: The use of an extra bit set to 0 or 1 as necessary to make the total number of 1 bits an even number; used as a means of error checking in data transmission. Compare **MARK parity**, **odd parity**, **space parity**.

event: A notification to an application of some occurrence, such as an interrupt created by a keypress, that the application may want to respond to.

event-driven: (1) Describes a kind of program that responds to user input in real time by repeatedly testing for events posted by interrupt routines. An event-driven program does nothing until it detects an event such as a click of the mouse button. (2) Describes a style of programming in which program actions are based on events generated by the user, rather than on some sort of fixed script.

Everyone: The user category to which you can assign privileges for any user with access to an AppleShare file server, whether logged on as a registered user or as a guest.

fiber optic cable: See **optical fiber**.

file: (1) Any named, ordered collection of information stored on a disk. Application programs and operating sys-

tems on disks are examples of files. You make a file when you create text or graphics, give the material a name, and save it to disk; in this sense, **file** is synonymous with **document**. A Macintosh file consists of a **data fork** and a **resource fork**. (2) For UNIX operating systems, an array of bytes; no other structure is implied by UNIX systems as they even treat devices like files.

file management: A general term for copying files, deleting files, and other chores involving the contents of disks.

file server: (1) A network device, usually consisting of a computer and one or more large capacity disks, on which network users can store files and applications in order to share them. (2) A specially equipped computer that allows network users to store and share information. (3) A combination of controller software and a mass-storage device that allows computer users to share common files and applications through a network. AppleShare software, Macintosh computers, and one or more hard disks make up a file server on an AppleTalk network system.

file transfer protocol: A protocol that exchanges files with a host computer.

Finder: The application that maintains the Macintosh desktop and starts up other programs at the request of the user. You use the Finder to manage documents and applications, and to get information to and from disks. You see the desktop upon starting up your computer, unless you have specified a different startup application.

Finder information: Information that the Finder provides to an application upon starting it, telling it which documents to open or print.

firmware: Software that has been written into nonchangeable memory, and thus does not need to be loaded into the system for each use. Most printers and output devices store their software in this form.



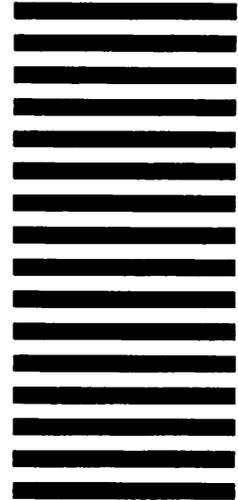
NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST CLASS MAIL PERMIT NO. 99 HOMEWOOD, IL

POSTAGE WILL BE PAID BY ADDRESSEE

BUSINESS ONE IRWIN
1818 Ridge Road
Homewood, IL 60430-9924



ATTENTION
Joan Zagone



BUSINESS ONE IRWIN is dedicated to providing you with the most timely, authoritative, and comprehensive computer guides to meet your personal and professional needs. Our goal is to offer you excellent reference tools written by experienced, highly-qualified authors. Please take a moment to answer the following questions so that we can continue to develop our products based on your needs.

Thank you.

BUSINESS ONE IRWIN

Current Computer Hardware:

- IBM or compatible
- Apple or Macintosh
- Commodore
- Atari

Disk Density:

- High
- Low

Disk Size:

- 5.25" 3.5"

Is It Portable?

- Yes No

Printer Type:

- LaserJet
- PostScript
- Dot Matrix
- Other

Peripherals:

- Modem
- Scanner
- Mouse
- Networking
- FAX
- Other (*please specify*)

Operating System:

- MS-DOS
- OS/2
- Macintosh
- UNIX

Programming Languages:

- BASIC
- C
- Pascal
- Prolog
- Assembly
- HyperTalk

The Most Recent Software Program

I Purchased Is: _____

Software I Use Most Often:

(please check all that apply and state name of product)

- Word Processing _____
- Spreadsheet _____
- Database/File Mgmt _____
- Desktop Publishing _____
- Presentation Graphics _____
- Accounting _____
- Statistical Analysis _____
- Financial/Tax Planning _____
- Integrated _____
- Business Form Design _____
- Utilities _____

Additional Comments: _____

Name _____

Title _____

Organization _____

Address _____

City/State/Zip _____

Telephone _____

- Please send a catalog of Business One Irwin books.

businessOne
IRWIN

1818 Ridge Road, Homewood, IL 60430

formatting: The process by which information is arranged into **packets**, or prepared for transmission on a network. Includes the addition of sender and destination addresses, and other transmission control information.

full duplex: A four-wire communication circuit or protocol that allows two-way data transmission between two points at the same time. Compare **half duplex**.

full-duplex communication: A method of data transmission where two devices transmit data simultaneously. This method allows the receiving device to echo back each character of your message as it is received.

full router: A router that connects two or more networks at the same site into an internetwork.

gateway: A device that connects networks that use different protocols. In effect, it translates between the protocols so that devices on the connected networks can exchange data.

GB: See **gigabyte**.

global backup: The process of backing up all the files on a hard disk. Compare **incremental backup**.

group: (1) A named collection of one or more registered file server users. Groups are created for users who usually have common interests and share information. (2) The user category to which you can assign access privileges for members of groups created by the AppleShare administrator. The administrator creates groups of registered users only.

group ID (gid): In A/UX, a number that indicates a group to which you belong at login time. As a member of a group, you have access to certain files and directories shared by other members of your group. Each user login name has at least one group ID associated with it.

guest: A user who is logged in to a server without a registered user name and password. A guest cannot own a private folder.

half duplex: A two-wire communication circuit or protocol designed for data transmission in either direction but not both directions simultaneously. Compare **full duplex**.

half-duplex communication: A way of communicating between your computer and another computer or a peripheral device in which you can only send data or receive it at one time—not both. The other computer cannot echo back each character of your message as it is received.

half-router: A router used with a modem that connects a user or entire network to a remote site. Two half-routers can create a wide area network.

handshaking: The exchange of status information between a DCE and a DTE used to control the transfer of data between them. The status information can be the state of a signal connecting the DCE and the DTE, or it can be in the form of a character transmitted with the rest of the data. See also **Data Carrier Detect**, **Data Set Ready**, **Data Terminal Ready**, **XOFF**, **XON**.

hardware handshake: A protocol that tells the computer to start or stop sending data by setting the DTR (Data Terminal Ready) line logic state. Also known as the Data Transfer Ready protocol. Compare **XON/XOFF**.

host computer: (1) A multi-user computer, such as a mini-computer or mainframe, that serves as a central processing unit for a number of terminals. (2) The computer that receives information from and sends data to terminals over telecommunication lines. The computer that is in control in a data communication network. The host computer may be a mainframe computer, minicomputer, or microcomputer.

incremental backup: The process of backing up all files on a hard disk that have been created, modified, or copied onto the disk since the last global backup.

information service: A large database that you can subscribe to for news, stock quotations, and other services and information. You communicate with these services through your computer equipped with a modem.

Installer: A utility program that lets you choose an Installation script for updating your system software or adding resources.

intelligent device: A device that contains a microprocessor and a program that allows the device to interpret data sent to it as commands that the device is to perform.

interactive protocol: A protocol that lets you communicate interactively with a host computer. In this kind of protocol, part or all of the contents of the screen memory are sent to the host when you press the Return key. You do not have to communicate with the host by sending it disk files.

interface: (n.) (1) The point at which independent systems or diverse groups interact. The devices, rules, or conventions by which one component of a system communicates with another. Also, the point of communication between a person and a computer. (2) The part of a program that defines constants, variables, and data structures, rather than procedures. In C, the compile-time and run-time linkage between your program and Toolbox routines. (3) The equipment that accepts electrical signals from one part of a computer system and renders them into a form that can be used by another part. (4) Hardware or software that links the computer to a device. (v.) To convert signals from one form to another and pass them between two pieces of equipment.

internet: (1) A network made up of two or more interconnected local area or wide area networks. (2) A worldwide, interconnected group of networks. Internally, the internet is

composed of heterogeneous networks (such as ARPANET and CSNET) that use different message formats and protocols. Through the use of gateways that convert formats and protocols between networks, the internet appears externally as a single network, with hosts on interconnected networks appearing as interconnected hosts. (3) Any interconnected group of networks, whether or not it's on the worldwide internet; for example, an interconnected group of AppleTalk network systems. Network users in an internet can share information and network devices.

internet address: (1) An address for a computer on a network. The internet address consists of a network number and a host number that is unique for that network. (2) The AppleTalk address and network number of a socket.

ISO/OSI model: See OSI model.

K: See kilobyte.

Kbit: See kilobit.

Kbyte: See kilobyte.

kHz: See kilohertz.

kilobit (Kbit): A unit of measurement, 1024 bits, commonly used in specifying the capacity of memory integrated circuits. Not to be confused with kilobyte.

kilobyte (K): A unit of measurement consisting of 1024 (2^{10}) bytes. Thus, 64K memory equals 65,536 bytes. The abbreviation K can also stand for the number 1024, in which case Kbyte or KB or is used for kilobyte. See also megabyte.

kilohertz (kHz): A unit of measurement of frequency, equal to 1000 hertz. See also megahertz.

LAN: See Local area network.

local area network (LAN): A group of computers connected for the purpose of sharing resources. The computers on a local area network are typically joined by a single transmission cable and are located within a small area such as a single building or section of a building. Compare **wide area network**.

local system: The computer from which a user originates a network command. Compare **remote system**.

local system administration: Management of a single computer. This includes such functions as starting up and shutting down the system, adding and removing user accounts, and backing up and restoring data. Compare **network administration**.

LocalTalk cable system: A system of cables, cable extenders, and connector boxes that connect computers and network devices as part of the AppleTalk network system.

LocalTalk connector: A piece of equipment consisting of a connection box, a short cable, and a machine-specific connector, that enables an Apple computer to be part of an AppleTalk network system. Also called a **LocalTalk connector box**.

login name: In UNIX systems, the name of a user's account. Used for identification purposes.

login prompt: The prompt (usually `login:` on UNIX systems) by which a system tells you it is ready to accept your login name.

log off: To indicate to a system or network that you have completed your work and are terminating interaction.

log on: To identify yourself to a system or network and start to use it. Usually logging on requires a password, depending on the system. Same as **log in**; opposite of **log off**.

mainframe computer: A central processing unit or computer that is larger and more powerful than a minicomputer or a personal computer (microcomputer). Frequently called **mainframe** for short.

MARK parity: A method of error checking in data transmission in which the most significant bit of every byte is set to 1. The receiving device checks for errors by looking for this value on each character. Compare **even parity**, **odd parity**, **space parity**.

media: The physical conductor of network transmissions, including electrical or optical fiber cables.

megabit (Mbit): A unit of measurement equal to 1,048,576 (2^{16}) bits, or 1024 kilobits, commonly used in specifying the capacity of memory ICs. Not to be confused with **megabyte**.

megabyte (MB): A unit of measurement equal to 1024 kilobytes, or 1,048,576 bytes. See also **kilobyte**.

megahertz (MHz): One million hertz. See also **kilohertz**.

MHz: See **megahertz**.

minicomputer: A multi-user computer, generally with more power than a personal computer yet not as large as a mainframe.

mode: (1) Any of several ways a computer interprets information. (2) A state of a computer or system that determines its behavior. A manner of operating.

modem: Short for **modulator/demodulator**; a peripheral device that links your computer to other computers and information services using the telephone lines.

modem command: An instruction to a computer system, usually typed from the keyboard, that directs a modem attached to the computer to perform some immediate action.

modem eliminator: A short cable for connecting two Data Terminal Equipment (DTE) devices together without a Data Communication Equipment (DCE) device.

MS-DOS: The Microsoft Disk Operating System. This is the operating system that governs the IBM PC (under the version PC-DOS) and compatible computers.

MultiFinder: A first-generation multitasking operating system for Macintosh computers that makes it possible to have several applications open at the same time, including background applications that let you perform one task while the computer performs another.

multilaunch application: An application stored on a server that several users can open and use at the same time.

multiplexing: Transmitting signals from multiple sources through a single medium.

multitasking: A process that allows a computer to perform two or more tasks during a given period of time; it is accomplished by alternating the actions of the computer between tasks. The method by which operating systems, such as A/UX, allow the user to open and run several applications at the same time. For example, multitasking would allow you to receive information from AppleLink, write a memo in MacWrite, and print an Excel spreadsheet—all at the same time.

multi-user: (adj.) Characterizes a mode or ability of an operating system to support several people using the same computer at once.

multi-user document: A document stored on a server that more than one user can open and make changes to at the same time.

multi-user system: An operating system, such as A/UX, that allows many users to access application software simultaneously.

name: The name presented to users of a network to identify a given network service.

naming protocol: A protocol used by AppleTalk to associate a name with the physical address of a network service.

network: A collection of interconnected, individually controlled computers, together with the hardware and software used to connect them. A network allows users to share data and peripheral devices such as printers and storage media, to exchange electronic mail, and so on.

network administration: Management of the software and hardware that connects computers in a network. This includes such functions as assigning addresses to hosts, maintaining network data files across the network, and setting up internetwork routing. Compare **local system administration**.

network administrator: The person who is responsible for setting up and maintaining the network.

network connection: A combination of hardware and software that lets you set up a particular implementation of the AppleTalk network system, such as LocalTalk or EtherTalk.

network device: A computer, printer, modem, terminal, or any other physical entity connected to a network.

network event: An event generated by the AppleTalk Manager.

Network File System (NFS): A protocol suite developed and licensed by Sun Microsystems that allows different makes of computers running different operating systems to easily share files and disk storage.

network manager: The person responsible for maintaining and troubleshooting the network.

network number: A unique number to each network in an internetwork that has been assigned by a seed router.

network system: A family of network components that work together because they observe compatible methods of communication.

NFS: See Network File System.

node: (1) A device that's attached to an AppleTalk network and communicates by means of the network. (2) Any network device that has an address on the network. (Some network devices, such as modems, may be connected to a network but not be a node themselves.)

node number: A number that distinguishes one node from all others on the network.

NuBus: An address bus and data bus incorporated into the system architecture of the Macintosh II. The NuBus architecture lets you add a variety of components to the system, by means of expansion cards installed in NuBus expansion slots inside the Macintosh II. NuBus is a trademark of Texas Instruments.

NuBus expansion slots: The six slots on the main circuit board of the Macintosh II to which you add cards for video monitors and peripheral devices, coprocessors, and network interfaces.

odd parity: The use of an extra bit in data transmission set to 0 or 1 as necessary to make the total number of 1 bits an odd number; used as a means of error checking. Compare even parity, MARK parity, space parity.

offline: (adj.) Not currently connected to or under the control of the computer. Used to refer to equipment such as printers and disk drives, information storage media such as disks, and the information they contain. Compare online.

online: (adj.) Currently connected to and under the control of the computer. Used to refer to equipment such as printers and disk drives, information storage media such as disks, and the information they contain. Compare **offline**.

open architecture: A computer system's ability to use a variety of optional components designed to meet specialized needs, such as video, coprocessing, networking, and so on. An "open" system is one to which a user with no technical background can easily add devices and expansion cards to customize the system.

optical fiber: A transmission medium that uses light to transmit a signal.

OSI model: The Open Systems Interconnection (OSI) reference model for describing network **protocols**, devised by the International Standards Organization (ISO); divides protocols into seven **layers** to standardize and simplify protocol definitions.

owner: (1) The registered AppleShare file server user who created a folder or was assigned ownership of a folder or volume. The owner is named in the Access Privileges window. (2) The AppleShare user category that the owners of folders or volumes use to assign access privileges to themselves.

packet: A unit of information that has been formatted for transmission on a network. See also **formatting**.

packet switching: A transmission method in which each data **packet** in a given transmission is sent independently. The sequence of packets traveling on a packet-switched network at any given time may originate from different senders and be headed for different destinations. Contrast **circuit switching**.

peripheral: A hardware device that is external to the essential circuitry in a computer. A hard disk and a printer are two examples.

Phase 1 AppleTalk: An AppleTalk protocol that can be used for nonextended networks.

Phase 2 AppleTalk: An AppleTalk protocol that makes extended networks possible.

port: (n.) (1) A socket on the back panel of a computer where you plug in a cable for connection to a network or a peripheral device. (2) A connection between the central processor unit and main memory or a device (such as a terminal) for transferring data. (v.) To move software from one hardware architecture to another.

primary group: The AppleShare group with whom you'll most often be sharing the documents you store on a server. Primary groups are specified by the administrator.

PrintMonitor: An application that monitors background printing and provides options intended to give you additional control over what happens to documents you are printing. See also **background processing**.

print queue: A collection of spooled documents, awaiting printing, stored on the print server disk and printed in order.

print server: A combination of software and hardware that stores documents sent to it over the AppleTalk network and manages the printing of those documents on a LaserWriter or other network printer.

print spooler: A utility that writes a representation of a document's printed image to disk or to memory, schedules it to print in a queue of other jobs, and then prints it.

protocols: The rules that govern interaction on a network. Protocols determine where, when, how, and in what format information is transmitted.

random-access memory (RAM): The part of the computer's memory that stores information temporarily while you're working on it. A computer with 512K RAM has 512 kilobytes

of memory available to the user. Information in RAM can be referred to in an arbitrary or random order, hence the term random-access. (As an analogy, a book is a random-access storage device in that it can be opened and read at any point.) RAM can contain both application programs and your own information. Information in RAM is temporary, gone forever if you switch the power off without saving it on a disk or other storage medium. An exception is the battery RAM, which stores settings such as the time and which is powered by a battery. (Technically, the read-only memory (ROM) is also random access, and what's called RAM should correctly be termed read-write memory.) Compare **read-only memory**.

read-only memory (ROM): Memory whose contents can be read but not changed; used for storing firmware. Information is placed into read-only memory once, during manufacture. It remains there permanently, even when the computer's power is turned off. Compare **random-access memory**.

remote: (adj.) At a distance. Unable to be connected directly using wires only, but requiring communications devices. Compare **local**.

remote computer: A computer other than your own but in communication with yours through telephone lines or other communication links. A remote computer can be at any distance from your computer, from right beside it to thousands of miles away.

remote system: On a network, any computer other than the local system.

remote site: A computer or network that is accessed through a long distance communications medium, such as telephone lines, ISDN, or a satellite.

repeater: A device that extends the maximum length of cable in a single network, so that the network can be expanded. Compare with **bridge**, **gateway**, and **router**.

resource: A file contained in the System Folder that provides information the microprocessor needs to communicate with certain devices attached to the computer system. A **printing resource** is a system file that lets you print on a corresponding printer attached to the computer.

ring topology: A layout scheme in which network devices are connected by the physical medium in a closed loop.

route: The path information takes when it is transmitted from one network device to another.

router: A device that connects similar networks to each other. A router receives data transmitted from other nodes and retransmits it to its proper destination over the most efficient route; this route may include several routers, each forwarding the data to the next. Compare with **bridge**, **gateway**, and **repeater**.

RS-232: A common standard for serial data communication interfaces.

RS-232 cable: Any cable that is wired in accordance with the RS-232 standard.

RS-422: A standard for serial data communication interfaces, different from the RS-232 standard in its electrical characteristics and in its use of differential pairs for data signals.

seed router: A router in that defines and contains the identifying information about a network. Every network in an internetwork must have at least one seed router.

server: A computer that provides a particular service across a network. The service may be file access, login access, file transfer, printing, and so on. Computers from which users initiate the service are called **clients**.

service: A specialized function that a network provides to users, such as file sharing and electronic mail.

shielded cable: A cable with a special metallic wrapping around its wires. This wrapping reduces radio-frequency interference.

SNA/SDLC: Abbreviation for **Systems Network Architecture/Synchronous Data Link Control**. SNA is a set of rules for controlling the transfer of information in a data communication network. SDLC is a protocol that uses commands to control data transfer over a communication line. IBM telecommunications products manufactured after 1978 use this protocol.

socket: On a network, a communication mechanism originally implemented on the BSD version of the UNIX operating system. Sockets are used as endpoints for sending and receiving data between computers.

space parity: A method of error checking in data transmission in which the most significant bit of every byte is set to 0. The receiving device checks for errors by looking for this value on each character. Compare even parity, MARK parity, odd parity.

spool: To send documents to the print server disk to be stored until they are printed. The word's origin is SPOOL, which stands for **simultaneous peripheral operations on line**.

spooler: See **print spooler**.

spool printing: (1) Writing a representation of a document's printed image to disk or to memory, and then printing it. (2) The printing of bitmapped text in which the print manager writes QuickDraw commands to a file or to memory; the printer driver then converts these commands into a bit image it sends as graphic data to the printer.

standard: A set of specifications for designing hardware or software that is recognized by multiple vendors, an official standards organization, or both.

star topology: A layout scheme in which network devices are arranged so that all are connected to a central controlling device.

synchronous: Able to perform two or more processes at the same time, such as sending and receiving data, by means of a mutual timing signal or clock. Compare **asynchronous**.

synchronous modem: A modem that provides two clocks for synchronous communication with its host computer: one clock for sending data from the host computer to the modem, and a second clock for sending data from the modem to the host computer.

synchronous transmission: A transmission process that uses a clocking signal to ensure an integral number of unit (time) intervals between any two characters. Compare **asynchronous transmission**.

TCP/IP: Abbreviation for **Transmission Control Protocol/Internet Protocol**; a suite of networking protocols developed at the University of California for the U.S. Department of Defense.

telecommunication: Transmitting information across varying distances, such as over telephone lines.

telecommunications: The science and technology of communication by electrical or electronic means.

terminal: A keyboard and display screen through which users can access a **host computer**.

terminal emulation: Software that enables a personal computer to communicate with a **host computer** by transmitting in the form used by the host's terminals.

token passing: An access method according to which devices on a network pass a special sequence of bits, known as the "token," from one device to the next. A device can only transmit data on the network if it is in possession of the token.

topology: The physical layout of a network.

traffic: Transmissions travelling across a network.

transeiver: A computer's hardware mechanism through which network transmissions are sent and received.

translate table: The table that controls character conversion between the Macintosh and the IBM host.

translation menu: A menu, such as "Mac to MS-DOS" or "Mac to Mac," that appears when two disks are shown in the Apple File Exchange window.

translator: The information that the Apple file exchange utility needs to translate a document created with an application on one operating system into a document that can be used with a similar application on another operating system.

translator file: A file containing one or more Apple File Exchange translators.

trunk topology: See **backbone topology**.

tty: A terminal; abbreviated from teletypewriter, which was the first terminal device used on UNIX operating systems.

twisted pair: A common, relatively low-cost network cable that consists of two insulated wires twisted about each other.

WAN: See **wide area network**.

wide area network: Computers and/or networks connected to each other using long distance communication methods,

such as telephone lines and satellites. Compare with **local area network**.

XOFF: A special character (value \$11) used for controlling the transfer of data between a DTE and a DCE. When one piece of equipment receives an XOFF character from the other, it stops transmitting characters until it receives an XON. See also **handshaking**, **XON**.

XON: A special character (value \$13) used for controlling the transfer of data between a DTE and a DCE. See also **handshaking**, **XOFF**.

XON/XOFF: A communications protocol that tells the computer to start or stop sending data by sending the appropriate character: either an XON or an XOFF. The ImageWriter LQ sends an XOFF when its input buffer is nearly full and an XON when it has room for more data. Compare **hardware handshake**.

Yellow Pages: A Network File System protocol for sharing a common database of user information across a local area network.

zone name: An identifier that allows groups of network devices on an AppleTalk internetwork to be displayed in separate lists within the Chooser.

Index

A

Alexandria Institute 334–336
AppleShare (Apple Computer Inc.) 138–140
AppleShare Print Server (Apple Computer Inc.) 160–161
Atkinson, Bill 58

B

Barlow, John Perry 338
Brand, Stewart 39, 328, 329, 330
Bressand, Albert 330
Burnham, David 323
Bush, Vannevar 19–22, 23
 “*As We May Think*” 20–22, 23
 Bush Rapid Selector 21
 hypertext 19
 links, associative 22
 memex 20–22
 Office of Scientific Research and Development 20
 thought, associative 22

C

collaboration
 brainstorming example 261–269
 computer conferencing example 250–253
 data backup example 269–276
 designing the physical space 290–291
 document management 229–230
 version control 229–230
 group authoring and editing 219–229
 information gathering 220
 organization analysis 220
 presentation analysis 220
 prewriting 219

- sequence analysis 220
- hypermedia support 296–297
- problem solving and 294–296
- project management example 254–261
- redefining meeting ecology 280–287
 - benefits of 286–287
 - consensus 285–286
 - role of facilitator 283
 - role of group memory 284–285
- removing inhibitions to 278–279
- shared space and 125–130
 - Macintosh emulation 129–130
 - Xerox Colab 126–130
 - Xerox Argnoter 128–130
 - Xerox Cognoter 127–128
 - brainstorming 127
 - evaluating 128
 - generating output 128
 - organizing 128
- virtual meeting example 246–250
- virtual meeting facilitation 211–219

Crick, Francis 58

cyberspace 338–340

D

DataClub (International Business Software Inc.) 150–155

Davis, Stanley 43, 44, 47

de Sola Pool, Ithiel 329

de Solla Price, Derek 58

DeKoven, Bernie 56, 286

Delphi method 198

Doerr, John 289

Doyle, Michael 279

E

Engelbart, Douglas 23–30, 122, 198, 303

- “A Conceptual Framework for the Augmentation of Man’s Intellect” 24, 27
- Augment 24–25, 303

- augmentation of intellect 23, 27–30
 - augmentation means 29–30
 - conceptual framework 28–30
 - definition of 27
- Bootstrap Institute 25
- Knowledge Workshop 25–26
- NLS 24–25, 122

F

- fiber optic 100–101
- fiber optic, see also ISDN

G

- Gibson, William 338
- GraceLAN (Technology Works) 187–195, 237–238, 239
- Greenwalt, Kent 323
- Grief, Irene 287
- groupware
 - approaches 117–122
 - calendar management 118
 - comprehensive workgroup support 122
 - computer conferencing 119–120
 - computer-based telephony extensions 117–118
 - computer-supported audio/video teleconferencing 120
 - computer-supported face-to-face meetings 119
 - conversational structuring 120–121
 - face-to-face meeting facilitation 117
 - group decision support 117
 - group memory management 121
 - group-authoring 119
 - non-human meeting participants 122
 - presentation support 118
 - project management 118
 - screen-sharing 119
 - spontaneous interaction 121
 - Xerox PARC's System Concepts Laboratory (SCL) 121
 - text filtering 120

- GoFer 120
- On Location 120
- definition of 19
- design issues 297–300
- managing information flow 287–289
- problems 313–336
 - interpersonal conflicts 314–319
 - social implications 319–336

H

- Hawken, Paul 44
- hypertext and hypermedia 301–312
 - characteristics of 303–307
 - design issues 309–312
 - modes of 307–308

I

- Illich, Ivan 123
- information
 - anxiety 48–52
 - immediacy 49–52
 - mass-customization of 47–48
- information economy 37–45
 - control within 38
 - information compression 40–41
 - information expansion 40
 - information float 42–43
 - information transportability 41
 - ownership 38
 - poverty and 39
- information theory 45–46
 - redundancy 45–46
- Innis, Harold 320
- Inspiration (Ceres Software) 261–269
- Integrated Services Digital Network, see ISDN
- ISDN (Integrated Services Digital Network) 94–99
 - Basic Rate Interface (BRI) 94–95
 - imaging 98
 - Primary Rate Interface (PRI) 95–96

resource sharing 98–99
screen-based telephony 97
ISDN (Integrated Services Digital Network), see also fiber
optic

J

Janis, Irving 59, 318
Jobs, Steve 55, 58, 337
Johansen, Robert 117

K

Kerr, Robert 335

L

language
 ability to shape thought 124
Larsen, Judith 326
Lennon, John 58
Library Awareness Program 331
LocalTalk Cable System (Apple Computer Inc.) 136

M

Mail*Link QM-MS (StarNine Technologies Inc.) 182
MarkUp (Mainstay Software) 220–229
McCartney, Paul 58
McDonnell Douglas 122
McLuhan, Marshall 320, 321, 338
Mead Data Central 332
MediaTracks (Farallon Computing Inc.) 322
meeting
 asynchronous 19
 redefining ecology of 280–287
 benefits 286–287
 consensus 285–286
 role of facilitator 283
 role of group memory 284–285
 synchronous 19
 types of 246–247

- virtual 211–219
- Metaphor Computer Systems 122
- Microsoft Mail (Microsoft Corporation) 176–181
- modem 82–86
 - data compression 83–85
 - error correction 83–85
 - modulation standards 85–86

N

- Naisbitt, John 325
- Nelson, Ted 30–35, 302, 333
 - super virtualities 34–35
 - thinkertoys 34–35
 - Xanadu 30–34
- NetUpdater (MDG Computer Services) 195–198
- network
 - access methods 80–81
 - architecture 68–73
 - communications protocols 68–70
 - ISO/OSI Model 69–70
 - topology 70–73
 - assessment 131–132
 - bus 70–71, 131
 - ring 71–72, 132
 - star 72–73, 132
 - hardware 64–67
 - cabling 65–67
 - assessment 132–133
 - coaxial 66, 134
 - optical fiber 67, 135
 - twisted-pair 65–66, 134, 135–137
 - peripherals 65
 - installing 236–238
 - internetwork 75–79
 - bridges 77–78
 - gateways 78–79
 - repeaters 75–76
 - routers 76–77
 - local area network (LAN) 73–74

- management 101–106, 184–198
 - accounting management 187
 - administration 104–106
 - assessing needs 101–104
 - configuration management 186
 - fault management 184–185
 - network poller 185
 - packet analyzer 185
 - traffic monitor 185
 - performance management 185
 - security management 186–187
- media 132–137
 - assessment 132–137
 - Ethernet 132
 - LocalTalk 132, 133–137
- mixed platform 106–116
 - Macintosh-to-NeXT 111–116
 - DataVizBridge 113–116
 - GatorBox 111
 - GatorShare 111
 - Macintosh-to-PC 109–110
 - MacLink Plus/PC 112
- nodes and addressing 79–80
- packets 80
- planning 232–236
- remote access 209–211
- security issues 242–245
 - access privileges 242–245
- services 86–92
 - electronic mail 91, 161–184
 - assessment 181–184
 - establishing 240–241
 - workgroup considerations 162–164
 - file service 86–91, 137–157
 - assessment 155–157
 - dedicated file service 87, 138–140
 - distributed file service 90–91, 150–155
 - establishing 238–240
 - nondedicated file service 88–90, 141–149

- System 7.0 88-90, 145-149
- printer service 92, 157-161
 - establishing 238-240
 - MultiFinder print spooling 157-160
- screen sharing 245-246
- software 68
- standards 93
- testing 241-245
- transmission methods 81
- wide area network (WAN) 74-75

P

- pairware 289
- Personal Server Network (Information Presentation Technologies) 141-145, 239
- PhoneNET Liaison (Farallon Computing Inc.) 187, 209-211
- PhoneNET System (Farallon Computing Inc.) 136-137, 236-238

Q

- QuickMail (CE Software Inc.) 164-176, 240-241, 250-253

R

- Retrospect (Dantz Development Corporation) 269-276
- Rogers, Everett 326
- Rozsak, Theodore 325

S

- Schrage, Michael 52, 53, 54, 56, 59, 61, 123, 124, 125, 126, 277, 288, 291, 317
- Shannon, Claude 45
- Straus, David 279

T

- teleconferencing 198-208
 - asynchronous teleconferencing 199
 - synchronous teleconferencing 199
- TeleFinder (Spider Island Software) 199-208, 253-254

Thermodynamics, Second Law of 45
Timbuktu (Farallon Computing Inc.) 129–130, 212–216,
245–250, 261–269
Timbuktu/Remote (Farallon Computing Inc.) 217–219
Toffler, Alvin 37, 42, 59, 289, 290, 326
Tognazzini, Bruce 46
TrafficWatch II (Farallon Computing Inc.) 185
Turoff, Murray 198

V

Vertelney, Harry 299
virtual reality 338–340

W

Watson, James 58
Winkler, Dan 58
workgroup
 brainstorming example 261–269
 collaboration 52–55, 57–61
 benefits 63
 conceptual collaboration 54
 formal collaboration 52
 informal collaboration 52
 modeling processes and ideas 61
 perception of how others add value 61
 role of language in 61
 technical collaboration 54–55
 computer conferencing example 250–253
 data backup example 269–276
 document management 229–230
 version control 229–230
 group authoring and editing 219–229
 information gathering 220
 organization analysis 220
 presentation analysis 220
 prewriting 219
 sequence analysis 220
 interpersonal computing 55–57
 project management example 254–261

typical example 130–131, 231–232
virtual meeting example 246–250
Wozniak, Steve 58
Wurman, Richard Saul 48, 50, 51

X

Xanadu 30–34, 302–303
publishing 32–33
Silver Stands 31
Xanadu Operating Company 30
xanalogical storage 31–32



0 0000 00157 088

WITHDRAWN
MITRE LIBRARY

The only book that shows you how to use Macintosh-specific groupware for different real-world tasks. Groupware and collaborative computing are emerging technologies that will prove to be the focal point for the personal computer industry in the 1990s. If the 1980s was the decade of personal computing, the 1990s will be the decade of interpersonal computing. It's important to lay a foundation for these concepts and to provide an overview for their implementation. This is the role that *Groupware for the Macintosh*[®] plays.

Author Michael Fraase offers ways to integrate groupware products into the networked Macintosh environment, including multiple-platform networks comprised of DOS machines, NeXT workstations, as well as Macintoshes. You'll find:

- Several step-by-step tutorials showing the use of Macintosh-specific groupware products in a real-world setting.
- An insider's perspective of group dynamics and how they relate to the success of groupware.
- Tips for avoiding the problems associated with this new set of technologies.
- Plus much more!

Author Profile . . .

Michael Fraase is the proprietor of Arts & Farces, a multi-faceted communications and professional services business specializing in hypermedia production, technical writing, desktop/electronic publishing, and software design. He is the author of *Farallon's MediaTracks*[™] (Business One Irwin, 1991), *Macintosh Hypermedia Volume I: Reference Guide* (Scott, Foresman, 1990), and *Macintosh Hypermedia Volume II: Uses and Implementations* (Scott, Foresman, 1990).

ISBN 1-55623-521-6



15-3513-01