Now covers
ac® II, SE, and Plus

The Apple® Macintosh® Book

THIRD EDITION

Cary Lu

Microsoft
PRESS
®

mac

The Apple® Macintosh® Book

THIRD EDITION

Mac

Revised by
Cary Lu
and Ellen W. Chu

Microsoft
PRESS®

PUBLISHED BY
Microsoft Press
A Division of Microsoft Corporation
16011 NE 36th Way, Box 97017, Redmond, Washington 98073-9717

Library of Congress Cataloging in Publication Data

Lu, Cary, 1945–
The Apple Macintosh book.
Updated ed. of: Mac. 1985.
Includes index.
1. Macintosh (Computer) I. Lu, Cary, 1945–
Mac. II. Title.
QA76.8.M3L8 1988 004.165 88-5114
ISBN 1-55615-110-1

Printed and bound in the United States of America.

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

Distributed to the book trade in the United States by Harper & Row.

Distributed to the book trade in Canada by General Publishing Company, Ltd.

Distributed to the book trade outside the United States and Canada by Penguin Books Ltd.

Penguin Books Ltd., Harmondsworth, Middlesex, England
Penguin Books Australia Ltd., Ringwood, Victoria, Australia
Penguin Books N.Z. Ltd., 182-190 Wairau Road, Auckland 10, New Zealand

British Cataloging in Publication Data available

Contents

Section Three

Section Four

Preface to the Third Edition

This book came out of a conversation in 1983 between Bill Gates, founder of Microsoft, and Steve Jobs, cofounder of Apple Computer. Jobs' development team was hard at work on the then-secret Macintosh, which was many months away from introduction. Microsoft was busy writing application programs for the Mac, and Gates suggested that Microsoft's new publishing division prepare a book as well.

Shortly afterwards, Nahum Stiskin, general manager of Microsoft Press at the time, asked Cary to write the book. He replied that he had never heard of any microcomputer interesting enough to write a whole book about. Nahum insisted that Cary go to the West Coast and see the Macintosh in action.

Here is the book.

Since it was introduced in 1984, the Apple Macintosh has been evolving. The "computer for the rest of us" is no longer a single microcomputer; the term *Macintosh* now applies to a family of models, including the Mac Plus, the Mac SE, and the Mac II. The newest models are more flexible and more sophisticated than the original, but they are also more complicated.

This book has also been evolving. Most of the introductory material found in the first and second editions has been shortened or eliminated to make room for discussing the increasingly diverse range of hardware and software offered to Macintosh users.

The flood of new Macintosh products has been good for the fortunes of Apple Computer, Inc., as well as for its customers. A great challenge now faces the company—the challenge of middle age. Can it continue to be innovative and responsive to its customers' needs? Can any one company meet the diverse demands of the microcomputer industry in the long run?

We'll cover that story in the fourth edition.

Cary Lu
Ellen W. Chu
April 1988
Seattle, Washington

Acknowledgments

Many people have helped in the preparation of this book. We thank those on the staff of Microsoft Press who gamely tackled this edition even after working on the last one, including Karen Meredith, Claudette Moore, Dave Rygmyr, and Lee Thomas. We also thank Darcie Furlan, Chris Kinata, Becky Johnson, Mary Ann Jones, and artist Rick Bourgoin. Even though she doesn't want to be mentioned, we are especially grateful to editorial director Patricia Combs.

One small person, Meredith Sarah An, made her own unique contribution: She arrived before her parents could complete the revisions and stayed long after the book went to press.

Cary Lu
Ellen W. Chu
April 1988
Seattle, Washington

Computers Should Work Like People

Computers are supposed to help you get work done quickly, easily, and effectively. So why have they become cloaked in mystique? Because most computers are difficult to use. So difficult, in fact, that you hear about "computer literacy" as if everybody must learn a new language. Computer enthusiasts haven't helped by talking computer jargon that obscures rather than clarifies. And so the mystique has grown: To work with a computer, you must think like a computer.

Nonsense. Computers should work the way people do.

The Apple Macintosh has changed people's minds about how computers should work. It is the first widely used computer that does not force you to change your language and work habits for its sake. Learning to use the Macintosh does require practice, but because most steps are analogous to the way people already work, learning is quick and easy.

The Macintosh is a visual computer, one operated as much with symbols and pictures as with words. The ideas behind it were originally developed at the Stanford Research Institute and the Xerox Palo Alto Research Center; in 1981 these ideas led to the landmark Xerox Star, the first commercial computer with a visual interface. The Apple Lisa, launched in 1983, brought the price of such machines below $10,000—much less than the Star, but still too high for most individuals and many businesses. Now the Macintosh offers a visual interface for considerably less, enabling many more people to enjoy its benefits.

In fact, the advantages of a visual interface are so compelling that all future microcomputers will have one. At present, some graphic Mac-style software has begun appearing for the IBM Personal Computer/AT and IBM's newest microcomputer line, the Personal System/2. Meanwhile, the internal design of Apple's latest Macintosh models—the Macintosh SE (system expansion) and Macintosh II (sometimes called the open Mac)—has become more flexible, thus broadening the range of people, tasks, and accessories that a Macintosh can accommodate.

Before the Macintosh, even people who liked computers accepted an initial period of suffering before they became productive or had any fun. Not so with the Macs. Whether you're new to computers or an old hand, Macintosh visual technology means those periods of frustration are over—or at least much shorter.

HOW TO USE THIS BOOK

This book discusses the Macintosh microcomputer family, highlighting the Macintosh SE and Macintosh II, both introduced in 1987. The book's four sections need not be read straight through. The opening section, "In the Beginning" (Chapters 1 and 2), introduces the Macintosh computers and describes basic operation and software. It suggests what to do if you have a Macintosh model older than the SE and compares the features of all six Macintosh models. If you already own a Mac, you can probably skip much of this first section.

The second section, "Understanding the Macintosh" (Chapters 3 through 8), takes up each part of the Macintosh in detail, explaining how each component works and interacts with the others.

The third section, "Working with Macintosh" (Chapters 9 through 17), examines specific types of software, desktop publishing, communications, problem solving, and other practical topics.

The last section, "For More Information" (Chapters 18 through 23), is aimed at users interested in more specialized subjects, including photography and advanced communication techniques, the AppleTalk network, and the future of microcomputers.

A technical appendix, glossary, and index complete the book.

In any extended comments about a product, we cite its version number. Many products will change in time, so some of these comments may no longer apply. Products that were substantially complete but not yet available for sale are described as *pre-release*. Products in the planning stage but not yet operational are noted in this way: "The Acme Company *plans* to release..." or "The Beta company has *announced*...."

Some products described in this book may not be available, others will have been superseded by better ones, and a few companies may have gone out of business by the time you read this. Use this book as a guide, and check with magazines and user groups for the latest information.

Sectio

In the Beginning

What is the Macintosh? How has it changed since it first appeared in 1984? Is any Macintosh model right for you? This section introduces you to the Macintosh family of microcomputers. It tells you how to set up a Macintosh and reviews major components and their functions.

1: The Macintosh Family

Will you save time using a computer? Maybe. In many situations, the time saved with a computer is surprisingly hard to measure. Often the total time you spend on a project remains much the same because any potential savings are lost in redoing and polishing your first effort. When you typed a memo on a typewriter, you probably didn't bother to retype it for small errors; you simply wrote a correction in the margin. With a computer, you will be tempted to edit the memo and reprint it—several times. The results are probably cleaner, but do you save time?

Sometimes a computer makes complex tasks too easy, and so you take less care. You can become so involved with processing words, for example, that you lose sight of what you are trying to say. You can fiddle with numbers so effortlessly that their underlying meaning can become obscured. If you had to do a financial projection by hand on paper, you might start more carefully, choose more realistic figures, do fewer experimental projections—and be done sooner.

For many tasks, though, computers do save both time and effort; performing many repetitive calculations or maintaining complex inventory lists is certainly best done with a computer. And some projects wouldn't even be possible without one.

Computers are made up of hardware and software. A computer's physical parts—keyboard, screen, printer, wires, and so on—constitute hardware. You can touch hardware.

Software, or programs, are the instructions you need to turn a collection of hardware into a word processor, a number cruncher, or a game machine. You can't touch software, although you can touch the hardware it is stored on. Hardware is useless without software, and vice versa.

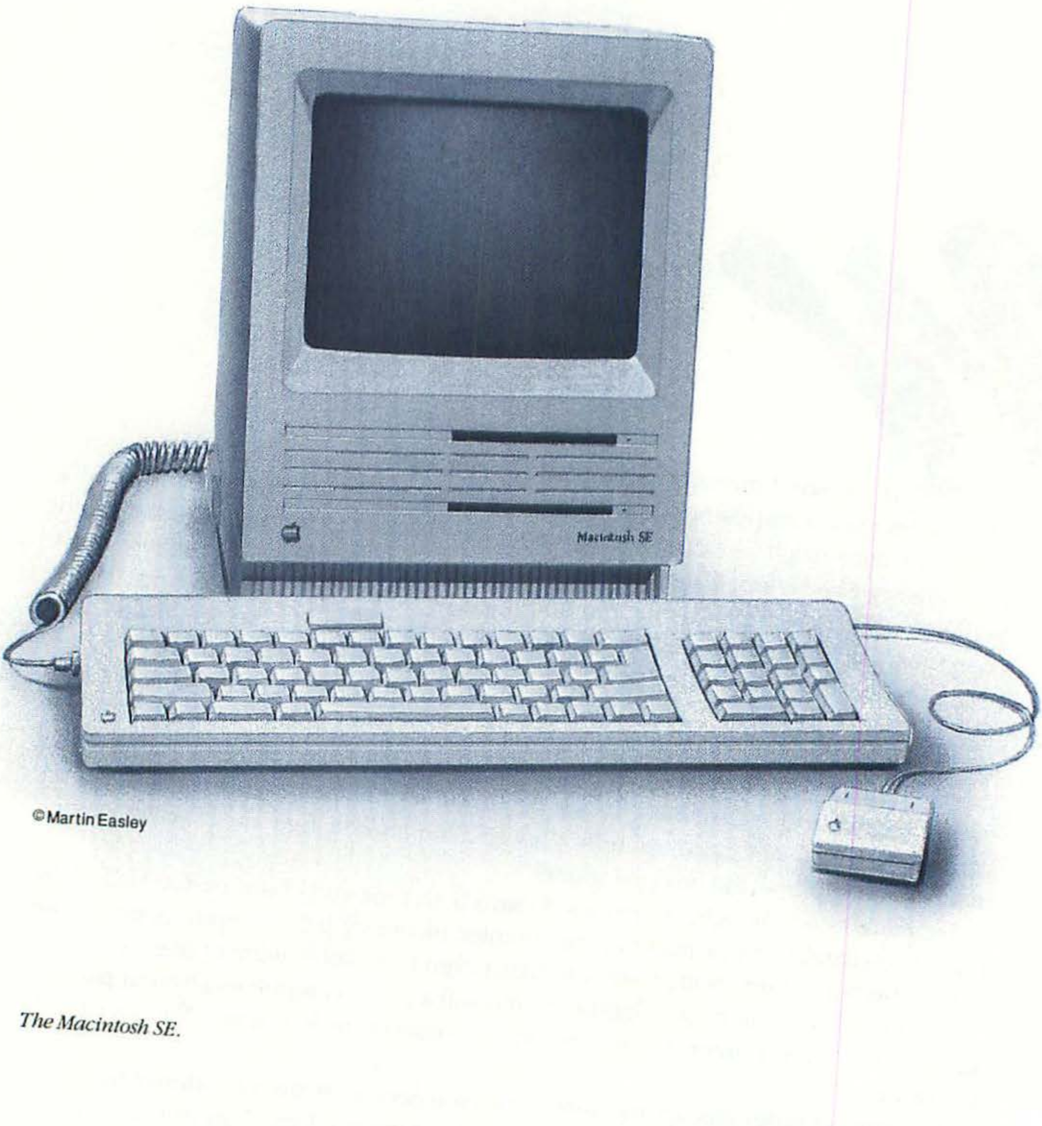

© Martin Easley

The Macintosh SE.

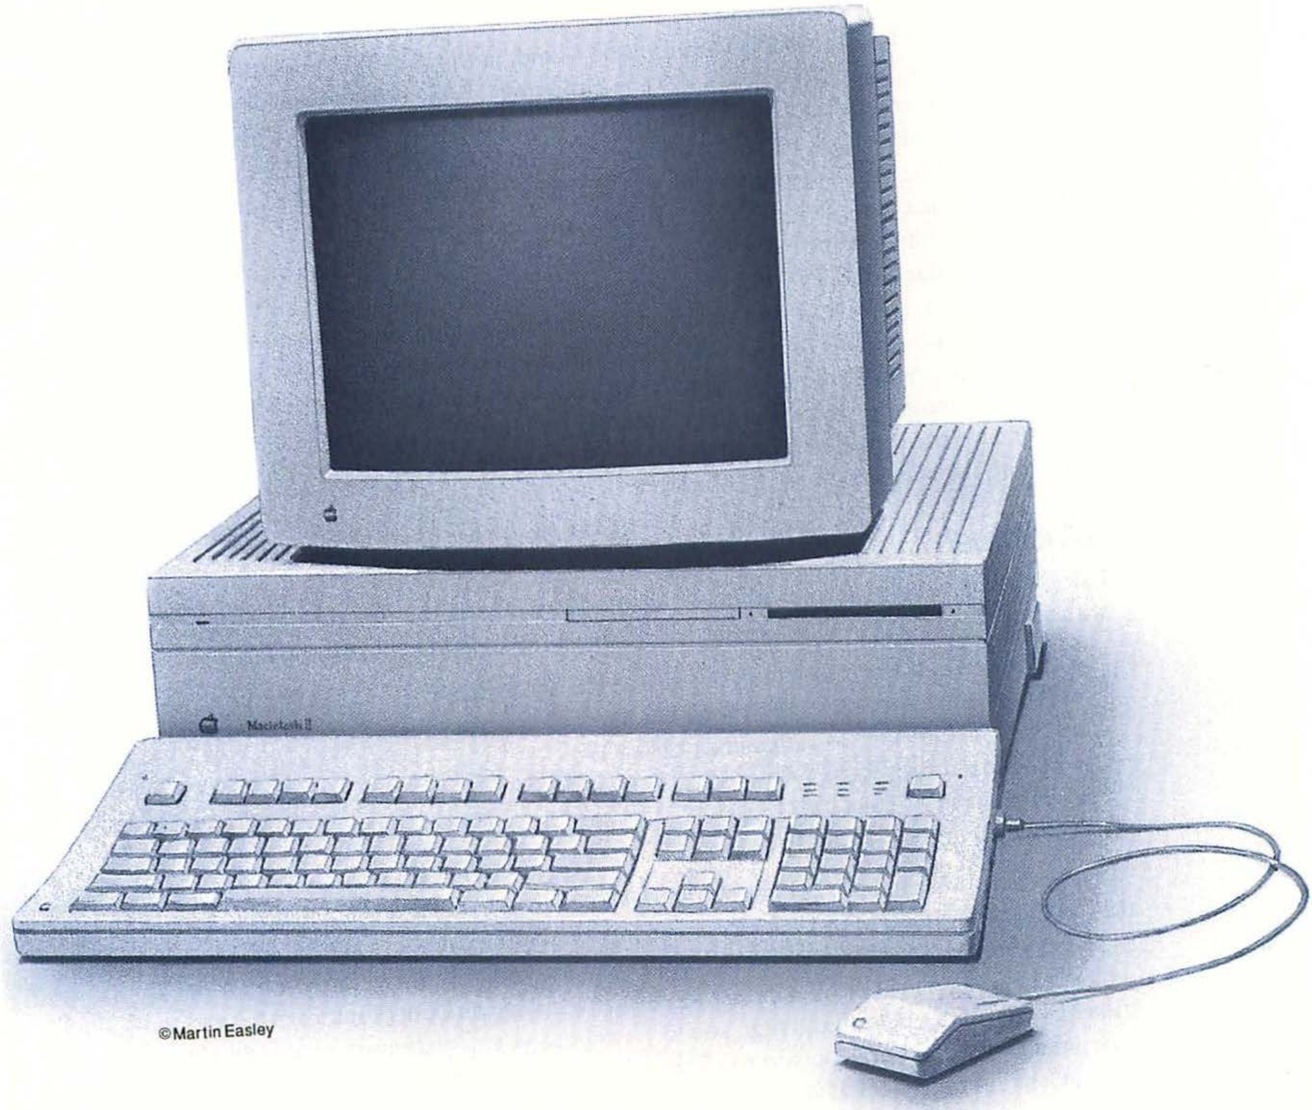

©Martin Easley

The Macintosh II.

IS A MACINTOSH RIGHT FOR YOU?

No computer can be all things to all people, and the Macintosh microcomputer line may not suit everyone.

The first Macintosh models were unique—and, at the time, incompatible with any other microcomputer. Now new software is making it easier for the Macintosh to work alongside or communicate with many other computers and devices. With accessory hardware, the Macintosh SE and Macintosh II can run software originally designed for the IBM PC. The Mac II handles shades of gray and color displays, which makes it suitable for complex computer-aided design and engineering projects.

This flexibility means that you could even use a Macintosh without a visual interface and run standard IBM PC programs that rely on cryptic commands rather than on graphics. Don't, however—not unless you are already using such programs and must continue to for a specific purpose, or unless you must work with people who have only traditional microcomputers. Not only is the Macintosh interface far easier to learn and use, but Macintosh programs are generally far better than traditional software. Today you can have the power and flexibility that have marked IBM PC-style computing right along with the accessibility of Macintosh visual technology.

Which Macintosh?

Using a Macintosh is easy; deciding which model to buy for the first time or what to do with your old Mac is not quite so easy. There have been six Macintoshes so far:

- The original Macintosh (1984) with 128 KB of memory
- The Macintosh 512K (1985) with 512 KB of memory
- The Macintosh 512K Enhanced (1986)
- The Macintosh Plus (1986)
- The Macintosh SE (1987)
- The Macintosh II (1987)

Only the last three models are currently being produced, and the Mac Plus will probably go out of production in 1988. For a comparison of all six models, see Table 1-1.

Base your choice of model on the kind of work you will do with it: Do you need color? A large display? Lots of memory for handling a large amount of information at one time? Does your work demand high speed, for repetitive calculations or processing long documents, for example? Do you work with information or programs that run only on IBM PC microcomputers? If you already own a Mac, are you tired of living with the compromises you might have made when you first bought it?

TABLE 1-1. A COMPARISON OF MACINTOSH MODELS

	Original Mac	*Mac 512K*	*Mac 512KE*	*Mac Plus*	*Mac SE*	*Mac II*
Processor	Motorola 68000	Motorola 68000	Motorola 68000	Motorola 68000	Motorola 68000	Motorola 68020
Processor speed	7.8336 MHz	7.8336 MHz	7.8336 MHz	7.8336 MHz	7.8336 MHz	15.6672 MHz
Memory (max.)	128 KB	512 KB	512 KB	1 MB (4 MB)	1 MB (4 MB)	1 MB (8 MB)
Expansion slots	None	None	None	None	1 for optional accessory cards	1 for video card 5 for optional accessory cards
Disk drives	1 internal floppy 1 external floppy	1 internal floppy 1 external floppy	1 internal floppy 1 external floppy	1 internal floppy 1 external floppy Optional external hard	2 internal floppy 1 external floppy 1 internal hard Optional external hard	3 internal: 1 internal floppy 1 internal hard Optional external hard 1 internal floppy or hard
Ports	2 serial Audio output Keyboard & mouse	2 serial Audio output Keyboard & mouse	2 serial Audio output Keyboard & mouse	2 serial SCSI Audio output Keyboard & mouse	2 serial SCSI Audio output Apple DeskTop Bus	2 serial SCSI Audio output Apple DeskTop Bus
Dimensions[1]	13.5x9.7x10.9	13.5x9.7x10.9	13.5x9.7x10.9	13.5x9.7x10.9	3.6x9.6x10.9	5.5x18.7x14.4 (main unit) 10x12.2x14.4 (b&w monitor)
Weight	16 lbs	16 lbs	16 lbs	16.5 lbs	17–21 lbs	24–26 lbs (main unit) 17–33 lbs (monitor)
System price (1988)	Discontinued	Discontinued	Discontinued	$1,500–$2,400	$1,800–$3,200	$4,000–$8,000

[1]In inches.

If you answered *Yes* to any of these questions, you may need to consider the new Macintosh SE or Macintosh II. Although you can add more memory, large displays, and speed-up accessories to the early Macs, doing so is clumsier than getting the right machine to begin with. On the other hand, both new Macs are noisier and more expensive than their predecessors (the Mac II is much bigger as well), so if desk space and cost are major issues, then you might want to stick with a Mac Plus or even an earlier model.

If you already own a Macintosh and are satisfied with your hardware and software, you may find little reason to change anything because of the new models. But realize that your hardware may not be able to accommodate the newest software. Many products now require the latest hardware to run; as their number grows, you will find yourself more and more at a disadvantage. If you choose not to upgrade, you will eventually be closed out of new developments, although repairs and replacement parts for all Macintosh models should be readily available for the next few years. Whatever your machine, if you buy any accessories, always try to buy those that are compatible with new models in case you decide to upgrade later.

WHAT YOU NEED

For an efficient Macintosh setup—whether based around a Mac SE, a Mac II, or an older model—here is a list of components you should consider. For many items, you can choose from several models and features; see the corresponding chapters in this book for more information.

The Essentials

- The main computer unit, including a mouse, one floppy disk drive, and system software.
- A keyboard (not included with the main Mac SE or Mac II computer unit).
- A box of blank microfloppy disks: either 3½-inch double-sided disks or 3½-inch high-density (1.6 MB) disks, if you have high-density floppy disk drives.
- A video card and video monitor (Mac II only).
- Some application software, such as a word-processing, database, or spreadsheet program.

The Nearly Essentials

- A hard disk drive, which makes computing easier by providing vastly more storage for programs and data and running much faster than floppy disk drives. Hard disk drives either plug in externally (Macs before the SE and II) or internally (Mac SE and II); internal hard disks are more convenient. On an SE, a 20 MB hard disk is sufficient for most personal computing; a 40 MB hard disk should take care of average business needs. On a Mac II, a 40 MB hard disk is the smallest practical size; for many applications, such as those with a great many graphics, an 80 MB or larger hard disk is essential.
- A second floppy disk drive, but only if you can't afford a hard disk drive or if you will be working in environments that might damage a fragile hard disk drive. Although having only a second floppy disk drive can be satisfactory with an SE, you should not consider buying a Mac II without a hard disk drive.
- At least two megabytes of random-access memory (RAM).

What You Will Probably Need

- A printer or access to one.
- Cables for connecting the computer to the printer: a serial cable for the ImageWriters, and AppleTalk (now called LocalTalk) cables if you are sharing ImageWriters on a network or for PostScript printers, such as all the LaserWriters except the IISC.
- A high-quality power extension cord.

What You Might Need

- Still more memory: You can easily add up to eight megabytes to the Mac II. As memory costs drop, adding a third and fourth megabyte to the SE will be affordable.
- A modem, which lets you communicate with other computers by telephone.
- Two modem cables—one connecting to the computer, the other to the telephone line; some modems come with these cables.
- Communication software.
- A mouse pad.
- Dustcovers.
- Carrying case (SE and other small Macs); the Mac II is awkward to carry, so you will probably leave it in one place.

TAKING CARE OF THE MACINTOSH

Although you should give your computer the same care and respect that you'd give any other valuable object, it is no more fragile than a television set. But do be sure, if you have a hard disk drive, not to jar the machine while the drive is running.

You cannot harm Macintosh hardware by any combination of typing on the keyboard or handling of the mouse. Generally, you can alter or erase software or lose data only by deliberate action. We'll go over some safekeeping steps as they come up.

TAKING CARE OF YOU

Don't work steadily at any computer for hours on end; plan to stretch your legs at least every hour or so. Don't forget to eat, and remember to talk occasionally with your fellow workers, family, and friends.

HOW MUCH DO YOU NEED TO LEARN?

Depending on your situation, you may need to learn only a single program—a word processor, perhaps. You may not even need to learn everything about that program. As you become more comfortable with the Mac, you'll probably want to learn more, but set your own pace. Learn only as much as you want, when you need to.

WHAT TO DO WITH YOUR OLD MAC

Now that new, expandable Macs are here, should you upgrade the Mac you have, or should you trade it in? That depends. If you don't need the latest and fastest hardware or software and your present Macintosh has enough memory (at least 512 KB), trading up is unnecessary. If, however, you need more speed or more features, or if you are shopping for your first Macintosh, buy a Mac SE or Mac II.

The original 128 KB Macintosh. If you have one of these, you should certainly upgrade its random-access memory (RAM) to 512 KB. Many independent companies, as well as Apple, can do this for you. Once the memory is upgraded, your machine will operate exactly like an original Macintosh 512K. Keep it, if the capacity, performance, and existing software are satisfactory for your needs. If you need any further enhancements, trade up to a new machine.

The original Macintosh 512K. The essential accessory for this model is an external disk drive. Used single-sided floppy disk drives should be readily available at low cost. If you need a hard disk drive, however, don't add one; trade in your machine for a newer model, and add the hard disk drive to the new machine. If this is impossible, then have a SCSI (pronounced SKUH-zee, for *small computer system interface*) port adapter installed and get a SCSI hard disk drive, preferably from the same vendor. You could also get an Apple HD 20 hard disk drive, which connects through the floppy disk port; older hard disk drives that connect through the serial ports may still be available, but they are unsatisfactory unless they are very cheap—under $150—and demonstrably reliable. SCSI connections are designed for speed, so a hard disk connected via a SCSI port will perform much better than one installed via either a floppy disk port or serial port.

If such an addition meets your needs, you may not have to switch to a more recent model. But the Mac 512K is already a dead end; an increasing number of new software products will not work with its read-only memory (ROM). Although this model can be partially upgraded to incorporate some features of the Macintosh Plus—such as the SCSI hard disk interface—the conversions are not really worth the trouble.

The Macintosh 512K Enhanced. This hybrid model combines some features of the Mac Plus with the basic layout of the Mac 512K. If you need more advanced features, you will do best by getting a new Mac. In other words, don't upgrade a 512K Enhanced; sell it instead. Apple built this model because dealers wanted a less expensive product to sell with the Macintosh Plus. Its unhappy design was outdated when it was introduced.

The Macintosh Plus. This was the first Mac that could support a hard disk drive adequately. If you have a Mac Plus, you should certainly consider getting an external hard disk drive. Internal memory upgrades up to four megabytes are also practical. Many other kinds of internal accessories—from hard disk drives to large display screens—exist for the Mac Plus, but internal additions are inadvisable now; in the long run, you will be better off switching to a Macintosh SE or II, which are both designed to accept a variety of internal components. On the other hand, if your work would gain nothing from such additions, you will not find much advantage to the slightly faster and noisier SE.

WRITING YOUR OWN PROGRAMS

If you find programming interesting and want to learn to program the Macintosh, by all means do so. But don't feel that you have to. Most people who use computers will never write a program. After all, excellent programs have already been written, so why duplicate the effort? Good word-processing software can take one skilled person several years to prepare; you may well prefer to do other things with that time.

As microcomputers become more common, chances are that a program meeting your specific needs will appear. But if you have a special requirement for which no program exists, you can write the program or have someone else write it for you. (Chapter 15 discusses programming languages.)

USING NEW PRODUCTS

Whenever you work with hardware or software for the first time, take a little extra care. Don't invest a lot of time entering data at the beginning; first be sure that the program does what you need. Enter a little information and check the results; print from the program to make sure everything works. Only after you are confident of both the program and your ability to use it should you invest time in intensive work.

KEEPING UP WITH NEW PRODUCTS

You may find yourself caught up in the craze to get the latest and greatest of each software or hardware accessory. If you enjoy doing this, fine; but remember, familiar products that work well may be enough. Glamorous features may be irrelevant to your needs; some striking new programs may be designed primarily for specialists. Remember, too, that every time you change programs, your work will inevitably be disrupted until you learn the new one. Although the Macintosh design makes this learning period short compared with learning time for other computers, you should weigh the potential benefits against the liabilities. In a business, be especially sure before changing programs and working procedures.

2: Getting Acquainted

Once you have decided on a Macintosh, you are ready to set it up and start working. Apple supplies an excellent manual with each Macintosh model, so this chapter presents mostly supplementary information.

FIRST STEPS

The first thing you need is a comfortable, properly lit place for your Mac. You probably want the keyboard and mouse at typing height, that is, a little lower than your desk or dining table. Adjust your chair so that you're seated comfortably; the fatigue that some people feel when working for long periods at a computer usually comes from poor lighting or an uncomfortable chair or desk setup. If two people of different heights share the computer, use an adjustable chair or keep a cushion handy.

Avoid placing your Macintosh next to a heating vent or radiator. Exposure to high temperatures might damage it or, more likely, its magnetic information-storage disks. Generally, if you can stand the temperature, so can a Mac.

Don't put the computer in direct sunlight, and be sure the area behind it is neither brilliantly lit nor very dark; otherwise, your eyes will have to readjust constantly for the difference. No lights should shine directly on the screen; position them to avoid glare. If you have bright overhead office lights, a shade or hood over the screen may help. (See Chapter 4 for more information about reducing glare.)

The only problem you are likely to have in assembling the computer is with the plugs. If you have trouble pushing a plug in, you have the wrong jack or the wrong orientation. Never force a plug! Apple's cables have small knobs to secure all plugs to the computer; other companies' cables often use small screws instead.

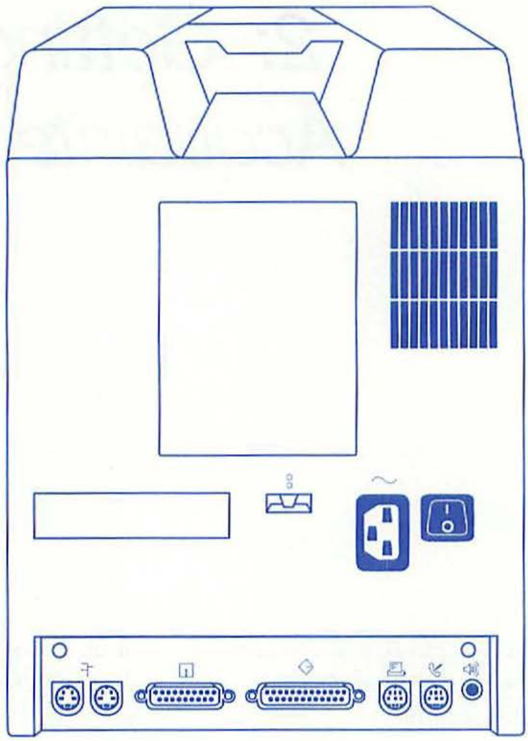

If the power supply is irregular in your area, or if you've had problems with other equipment, see Chapter 17 for suggestions. Outside North America, be sure that the electrical supply is the correct voltage. The power-line frequency (50 or 60 Hz) doesn't matter because all Macs can run on either frequency, but the voltage must be correct. Any transformer you use should have a capacity of 150 watts, or 150 VA, 350 watts for the Mac II.

A little plastic piece labeled INTERRUPT RESET, which allows you to restart the Mac without using the power switch, also comes with your computer unit. It fits on the left side of all Macs through the SE (the small Macs) and on the right side of the Mac II. You won't need to install this piece for now, so you can put it away.

If you ever need to move your Macintosh, the safest way is to disconnect all the cables and move each unit separately. Because they are fragile, hard disk drives should be moved with special care.

A FEW WORDS ABOUT EACH COMPONENT

Here's a quick rundown on each component. You won't need more information to begin operating a Macintosh successfully; for more detail, see Chapter 3.

The Main Computer Unit

The main unit of all models except the Mac II includes the display screen, or monitor, so you will have to put the entire computer on your work table. The Mac II's display screen comes separately. It can be set atop the main Mac II unit or placed directly on your desk or on a low monitor shelf resting on or above your desk. Position the monitor so that you can look straight at it without stooping or craning your neck.

Because the Mac II fan is noisy, you might want to place the main unit off to the side of your work space; be sure the monitor, mouse, and keyboard cables are long enough for you to do this. (You can get longer cables or extenders from your dealer, or you can make your own.)

Check the manual for how to turn on your main Mac unit. All the small Macs have a power switch on the rear left side; the Mac II turns on from the keyboard. To adjust screen brightness, use the brightness knob in front of or to the side of the display screen. You can leave the machine on constantly if you want; simply turn the brightness down whenever you're not going to use it for long periods. Alternatively, you can install a utility program that blanks the screen automatically after a period of inactivity; moving the mouse restores the screen.

Disks and Disk Drives

Disks, which store information magnetically, can be either floppy or hard. Every Macintosh comes with at least three start-up microfloppy disks; in addition, the Mac SE and Mac II may contain an internal hard disk—a rigid metal or glass platter with a magnetic coating. (See Chapter 6.) Hard disks can store a great deal more information than floppy disks; they also give faster access to that information.

The floppy disks fit, one at a time, into the disk drive slot(s) just below the display screen or, on the Mac II, on the right side of the main unit. They will only go in right side up; the plastic disk jacket has a small arrow showing the correct orientation.

Disks store several kinds of information in units called files. Functionally, a computer's files resemble the paper file folders in your filing cabinet. Stored information includes:

- *System and Finder information:* When you first turn on the Macintosh, you must insert a start-up disk in the disk drive. The computer transfers System and Finder information, which controls basic operation, from the disk into memory and onto the screen.
- *Application programs:* In many cases, programs such as MacPaint and MacWrite, which turn your Macintosh into a graphics tool or word processor, will be on the same disk as the System and Finder. Once the Finder is in memory, you'll need to indicate the program you want, and it too will be transferred into memory for active use.

- *Data:* Information you enter when using a program and other information needed by a program is stored in data files. Data files don't always have to be on the same disk with the program, although some programs will be easier to use if these files are handy. You can keep these different kinds of information on the same disk, but if they must share space with application programs as well as the System and Finder, you may run out of storage space on a microfloppy disk; a hard disk can alleviate this problem.

The Keyboard

The keyboard (and mouse) enables you to enter information and tell the Macintosh what to do with it. At least four Apple keyboards exist for the Macintosh:

- The original Macintosh keyboard, with or without a separate numeric keypad
- The Mac Plus keyboard, with built-in numeric keypad and arrow (cursor) keys
- The Apple Keyboard for the Mac SE and Mac II, with built-in numeric keypad and arrow keys
- The Apple Extended Keyboard for the Mac SE and Mac II, with built-in numeric keypad, arrow keys, and special function keys

Other manufacturers are also starting to produce keyboards that work with the Macintosh. Like the Apple Extended Keyboard, many of these copy the layout of the IBM PC keyboards.

All Macintosh keyboards have the standard typewriter keys plus some extra ones. The typewriter-like keys produce familiar letters and symbols but, unlike most typewriter keys, will repeat their characters automatically if you hold them down.

The Macintosh Shift key works like a typewriter's: When you hold it down, you get uppercase letters and punctuation marks instead of lowercase letters and numbers. But the Caps Lock key works a little differently than a typewriter's shift lock: It stays engaged when you press it and doesn't release until you press it a second time. When it is engaged, Caps Lock gives capital letters, as you would expect, but none of the symbols above the number or punctuation keys. For these you must use the Shift key even when Caps Lock is engaged.

The backspace key, labeled Delete on the Apple and Apple Extended keyboards, not only backs up as on a typewriter, but erases the letter(s) you have just typed.

Two keys peculiar to Macintosh keyboards, Option and Command (⌘, called the Apple key on the Mac SE and Mac II), work like two more shift keys, functioning only in combination with another key. The Option key serves several purposes, including generating special symbols and graphic elements. The Command key is a shortcut for telling the computer to do something such as eject a disk or delete a word.

The cursor keys on the numeric keypad and Mac Plus, SE, and II keyboards are used in some programs for moving the flashing insertion point, or cursor, on the screen, thus changing the location where your typing appears.

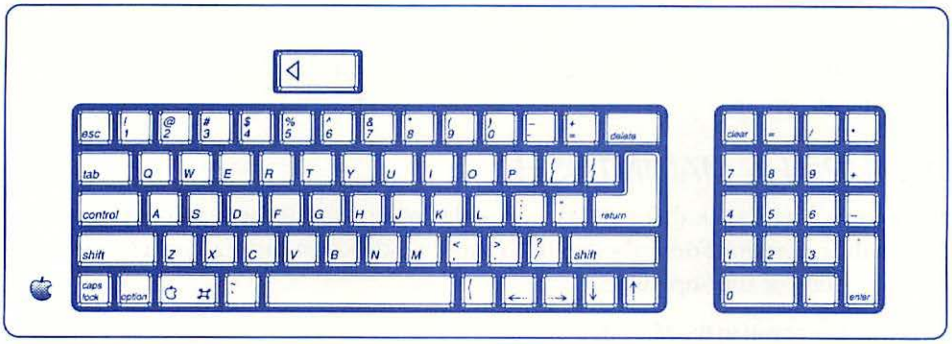

The Apple Keyboard.

The Apple Extended Keyboard.

The Enter key works differently from program to program, sometimes duplicating the results of pressing the Return key, sometimes simply telling your Macintosh to accept some typing or to execute a command. The Esc, Clear, and ◁ keys—found on Mac SE and II keyboards—have a variety of functions that may differ from program to program.

Finally, the Apple Extended Keyboard and some keyboards from other manufacturers have about 15 special keys called function keys, which can represent specific characters or perform specific operations. On the Mac SE and Mac II, these keys work with software written for the IBM PC as well as software designed for the Mac.

A word of warning: If you are in the habit of typing the letter *l* for the number *1*, or a letter *O* for zero, you'll have to change your ways. Sorry. The difference is very important to computers.

The Mouse

Along with the keyboard, the mouse lets you position your typing and point to items on the Macintosh screen.

STARTING UP THE MACINTOSH

Any floppy or hard disk that will start up a Macintosh has all the features you need to perform the operations described in this chapter. A disk that can start up a Macintosh is called a system, or start-up, disk.

Turn on the power to the Macintosh.

On a Mac II, pushing the ◁ key turns on the power and starts up the machine. If the Mac you are starting up has an external hard disk,

Turn on the hard disk, and then turn on the Mac.

If you are starting up from a hard disk and nothing happens, you may first have to initialize, or format, the disk; see the manual that accompanies the hard disk for how to do this.

If the Mac you are starting up has only floppy disk drives,

Insert a system disk into the disk drive slot.

In all cases, the disk drive whirs as it loads start-up information from the disk into the computer's electronic memory. A "happy Mac" icon appears on the screen.

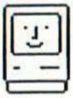

LOOKING AT THE DESKTOP

After the disk drive stops, you see the Desktop on the screen; this is part of the Macintosh's visual interface. The icons, or pictures, on the Desktop represent places for finding, storing, and discarding information. The first icons you see represent:

- The disk you've just inserted or your hard disk
- A trash can

A window may open as well. If so, try to ignore it for the next five minutes. If it gets in the way, scan ahead to the section called "Closing a Window."

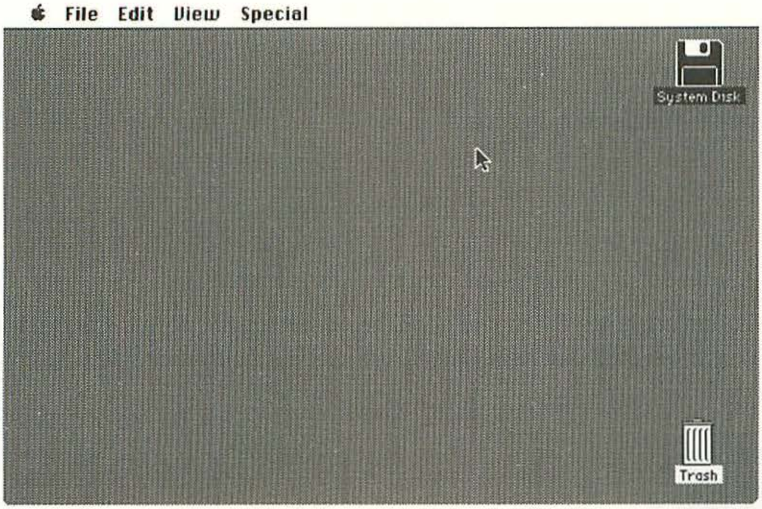

USING THE MOUSE

Move the mouse around your desk.

The pointer on the screen follows the mouse movements. You cannot move the pointer off the screen.

The pointer you see now is shaped like an arrow, but it can take on other shapes, as you'll see later.

Pointing

Move the pointer over an icon.

The exact spot you are pointing to depends on the pointer's shape. When the pointer is an arrow, position the arrow tip over the object. (Your disk icon may look different from this illustration.)

Clicking

Point at the Trash icon.

Press and release the mouse button once.

This action is called clicking. When you click the white Trash icon, it turns black, indicating that you have selected that icon. Selecting an icon means your next action will apply to that icon.

Move the mouse to the disk icon and click it.

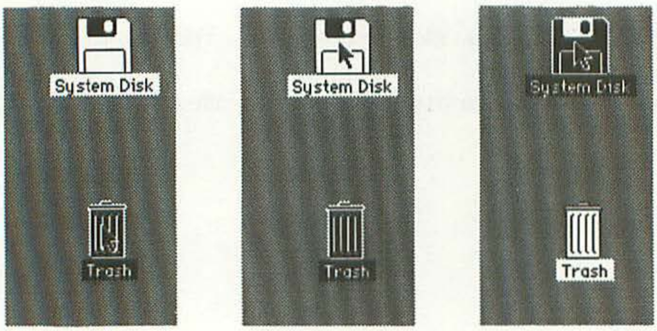

The Trash icon turns white again, and the disk icon turns black.

Click anywhere on the Desktop outside the icons.

This cancels your previous selection. Now none of the icons on the Desktop is selected.

Dragging

Move the pointer over the Trash icon.

Press and hold down the mouse button.

Move the mouse while holding down the button.

You use this essential operation, called dragging, in many ways. As you move the mouse, the icon's outline moves with it.

Release the mouse button.

When you release the button, the icon pops over to the location of its outline.

Pressing

Pressing means that you position the pointer and then press the mouse button, holding it until an action is complete. Pressing differs from dragging in that you don't move the mouse while holding down the button.

OPENING AN ICON

To find out what information an icon represents, you must open the icon; you can think of this as opening a drawer in a filing cabinet.

Click the disk icon (the system disk in this example).

Move the pointer to the menu bar at the top of the screen.

Point at the word File.

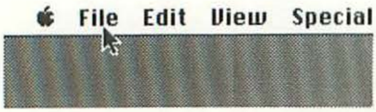

Press and hold down the mouse button without moving the mouse.

You've "pulled down" the File menu. The menu entries show what you can do with the icon you've selected—in this case, the system disk. Available menu choices appear in black letters; unavailable ones, such as Close, appear in dimmed or gray letters.

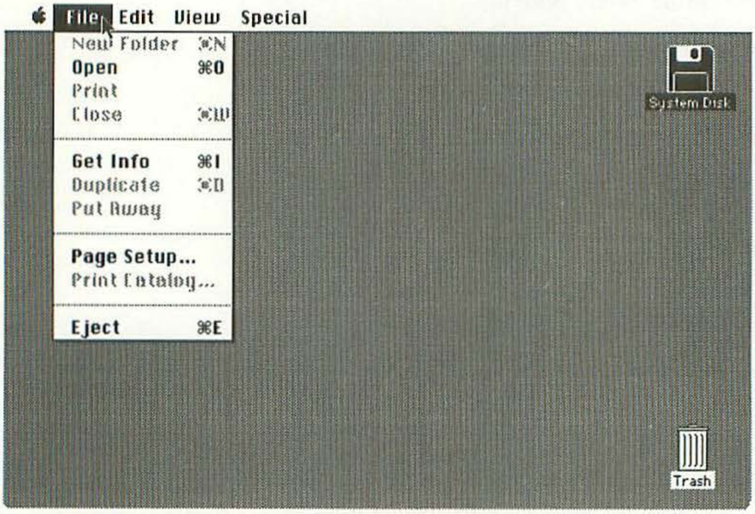

Still holding down the mouse button, move the pointer down the menu.

The available choices are highlighted in turn.

Position the pointer over Open.

The word Open is highlighted.

Release the mouse button.

You choose a menu item by releasing the mouse button when the item is highlighted. In Macintosh terms, you have chosen Open from the File menu.

The disk window now opens up.

Double-Clicking

To open an icon faster, point at it and quickly click the mouse button twice; this action is called double-clicking.

LOOKING AT THE DISK WINDOW

Each icon in a disk window represents a file on that disk. Some files can be application programs; others can be documents, graphs, or drawings; still others contain software the Macintosh itself needs to operate or resources useful when you are working with the application programs. The operations described here apply to most Macintosh windows, not only the disk window.

Moving a Window

You may need to reorganize your Desktop by moving windows around on it.

Point at the lined title bar at the top of the window, and drag it with the mouse.

Dragging a window's title bar moves the entire window.

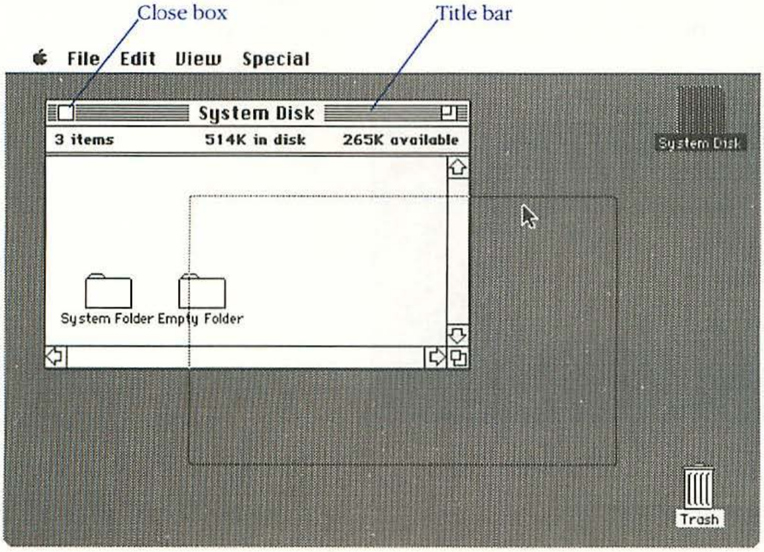

Release the mouse button.

The window moves to the location of its outline.

Changing a Window's Size

To reduce clutter on your Desktop, you can make windows smaller; if you want to see more of a disk's contents, you can make its window larger.

Point at the size box in the lower right corner of the window, and drag it with the mouse.

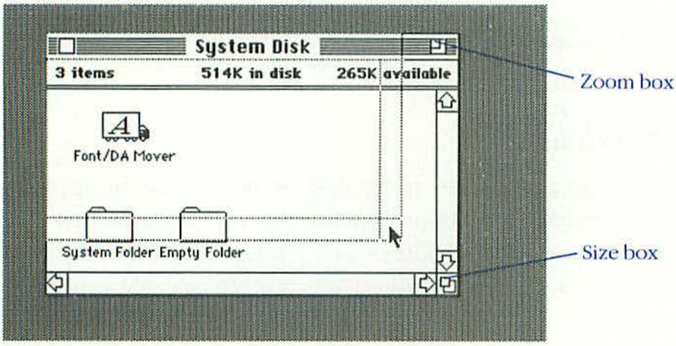

Notice that some of the icons drop out of sight as the window becomes smaller. You can also resize the window by clicking the zoom box. Try it; the window alternates between the maximum size possible on your Desktop and the size you've set.

Scrolling to See More of a Window

Sometimes a disk may contain more files, and therefore more icons, than you can see in the window at one time; to move the invisible icons into view, you need to scroll them past the window.

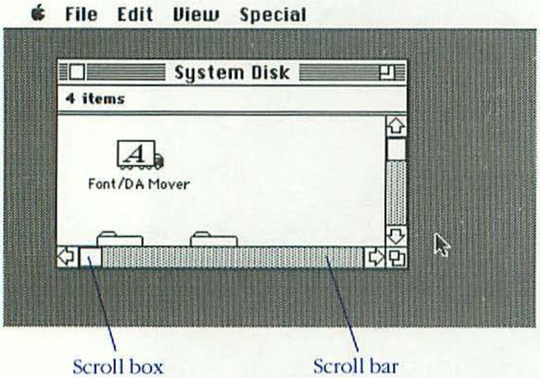

Click one of the scroll arrows.

Clicking the arrows moves the disk's contents past the window by a small increment (one line if you're looking at text).

Click in the gray area of one of the scroll bars.

Clicking in the gray area instead of clicking the scroll arrows moves the disk's contents past the window by a larger increment (a screenful if you're looking at text).

Position the pointer over a scroll arrow, and press the mouse button without moving the mouse.

Pressing a scroll arrow makes the disk contents scroll past the window until you release the button or until there is nothing more to scroll.

Drag the scroll box.

This action also brings hidden icons into view. The scroll box indicates the relative position of a window with respect to the contents of a disk or document.

Some windows do not permit resizing or scrolling and therefore lack a size box or one or both scroll bars. If a disk window is big enough to display all the contents of a disk, the scroll bars turn solid white, and scrolling is impossible.

OPENING MORE THAN ONE WINDOW

You can have several windows open on the screen at the same time.

Double-click the Trash icon.

(Remember, point at the icon, and quickly click twice on the mouse button.) This opens another window. It's empty because there is nothing in the Trash.

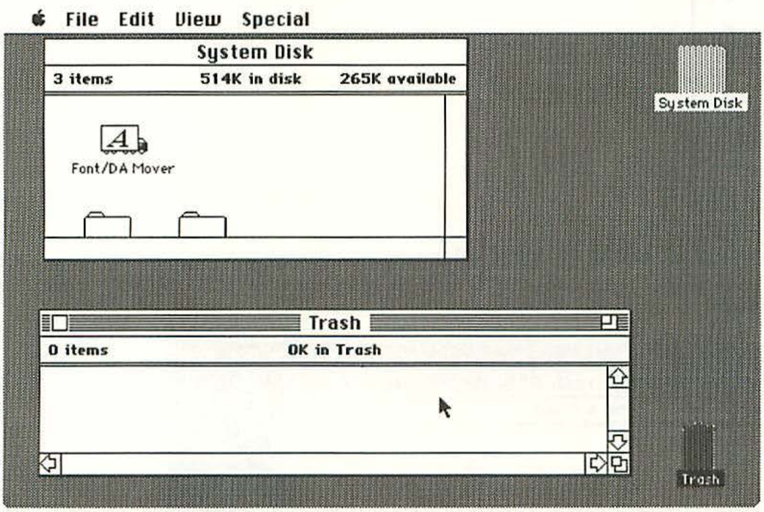

When more than one window is open on the Desktop at once, only one is active. The active window appears in front of, and may obscure, all the others; its title bar has horizontal lines on it. Notice that the lines have disappeared from the title bar of the disk window because it is no longer active. (Move or resize the Trash window if it has covered up the disk window.)

Click anywhere in the disk window.

The horizontal lines reappear in the disk window's title bar, showing that the disk window is active again.

Drag the size box of the disk window so that the window overlaps but does not obscure the Trash window.

Click anywhere in the Trash window.

Notice how the active window always shifts in front of any others.

Closing a Window

You can close your windows in either of two ways:

Choose Close from the File menu; *or*

Click the window's close box in its upper left corner.

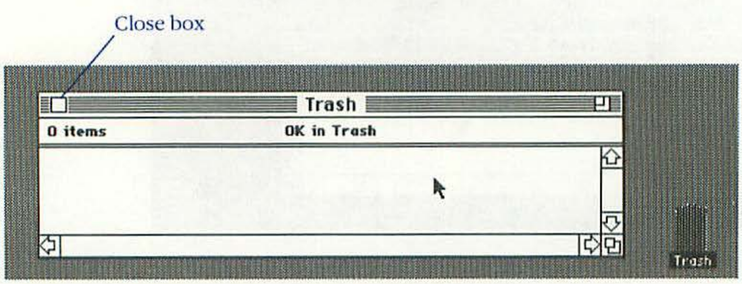

Close the Trash window; leave the disk window open.

WHAT THE DISK WINDOW TELLS YOU

Now that you can manipulate windows, let's take a closer look at the system disk window.

Click the disk window to make it active.

Drag its size box to make a large window.

The words and icons you see tell you which items you have stored on the system disk and how much space they take up.

A disk can hold only a finite amount of information. The top line in a disk window tells you how many items are on the disk, the amount of space taken up by the files on that disk, and how much space is still available for additional files. The units are thousands of characters, or kilobytes (abbreviated K in the window).

You can use the menus to find out more information about the files on the disk.

Choose by Name from the View menu.

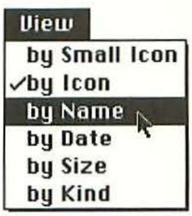

Point at View on the menu bar, drag the mouse to highlight by Name, and release the button. The icons disappear, and a list of files appears in alphabetical order. Along with each file's name, you see what size and kind it is and the date the file was last modified.

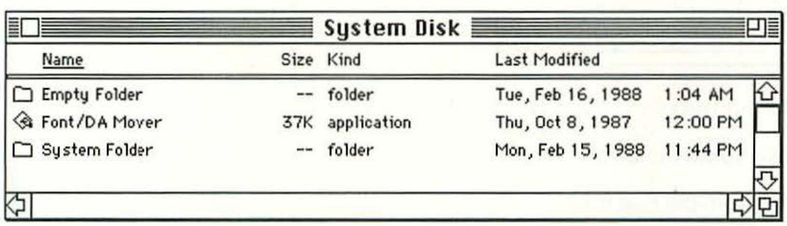

Pull down the View menu again, holding down the mouse button.

These choices represent ways you can organize the list of files. If you choose by Date, for example, the files are listed starting with the one most recently changed.

Drag to by Icon and release the mouse button.

Click any file icon.

Choose Get Info from the File menu.

A new window opens, showing information about the file.

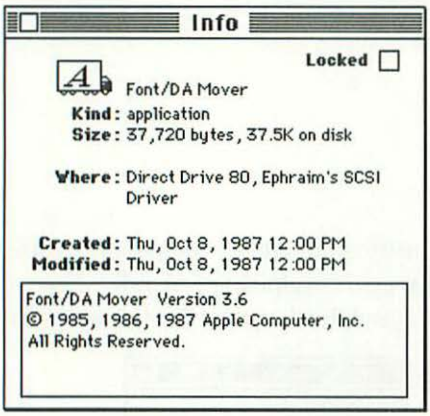

To close the Get Info window, you can either:

Click the close box in the Get Info window; *or*

Choose Close from the File menu.

TURNING OFF A MAC

If you are going to take a short break—up to a few hours—you can leave the Mac running. If you are going away for a longer period and want to turn off the Mac, don't simply shut off the power.

Choose Close from the File menu *and*

Choose Eject from the File menu; *or*

Choose Shut Down from the Special menu.

If you choose Shut Down, the Macintosh adjusts the hard disk drive so that it can be safely shut down, and, after a few seconds, it pushes any floppy disk out of the disk drive. Now you can turn off the power.

To restart, simply turn on the power and repeat the start-up procedure.

DESK ACCESSORIES

The Apple menu in the menu bar opens the way to the Macintosh desk accessories. You can use desk accessories at any time, no matter what else you are doing—whether or not you are using a program or whether or not any windows are open.

Point at the Apple symbol at the far left of the menu bar.

Press and hold the mouse button.

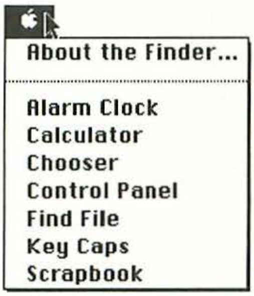

Drag the mouse down to highlight About the Finder…

Release the button.

Some information about the important Finder program appears on the screen. The Finder creates your Desktop, shows you which files are on the disk, copies files from disk to disk, and much more. You've been using the Finder throughout this chapter.

Click anywhere to get rid of the Finder information.

Go back to the Apple symbol and pull down its menu again.

The rest of the items in the menu are desk accessories. (Your menu may not be the same as this one.)

Choose each accessory in turn.

Click the accessory's close box to put it away.

What Do the Accessories Do?

Scrapbook saves material from a program and lets you move that information within the same program or to another program.

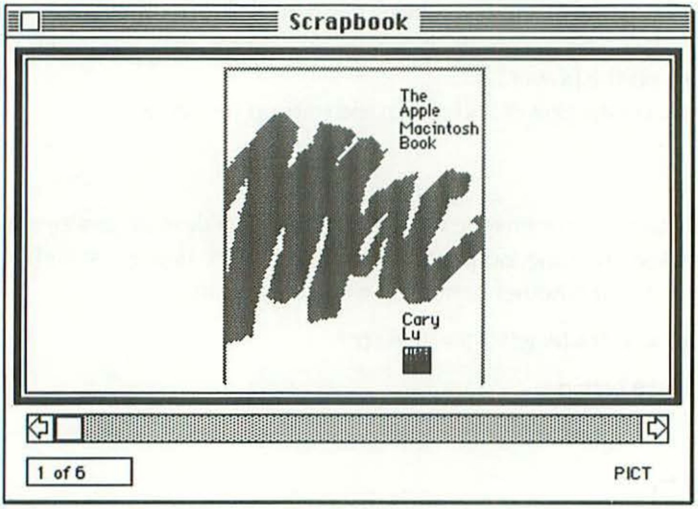

Alarm Clock displays a small box with the current time. Click the lever on the right side to expand the display.

Note Pad gives you a place to type text notes. You can turn pages (the Note Pad gives you eight) by clicking the upturned corner. Clicking the lower left corner below the diagonal line flips the Note Pad back a page. The Mac beeps if you try to put too much text on a sheet. This accessory does not always come with a Macintosh.

Key Caps displays a miniature Macintosh keyboard that tells you which characters your real keyboard generates. Press a few keys on the real keyboard or click them on the Key Caps display with the mouse and watch what happens; the keys are highlighted on the Key Caps display, and what they type appears in the blank bar above it.

Hold down Shift, Option, Caps Lock, or a combination of these keys on the real keyboard, and watch what happens; the display keyboard is instantly transformed into a variety of symbols and graphics characters. To see what's available in a different typestyle (font), pull down the Key Caps menu and select the font you want.

Chooser lets you select a printer type and, if you are on a network, a specific printer. It also lets you connect or disconnect the AppleTalk network.

Control Panel lets you adjust a variety of parameters. The general settings include the volume of the Macintosh speaker (drag the sliding control up or down), the Desktop pattern, and the size of the RAM cache (see Chapter 14). For other settings, click the icons that appear on the left side; for example, clicking the keyboard will bring up keyboard settings such as how fast the keyboard keys repeat (turtle for slow, hare for fast). See your manual for details.

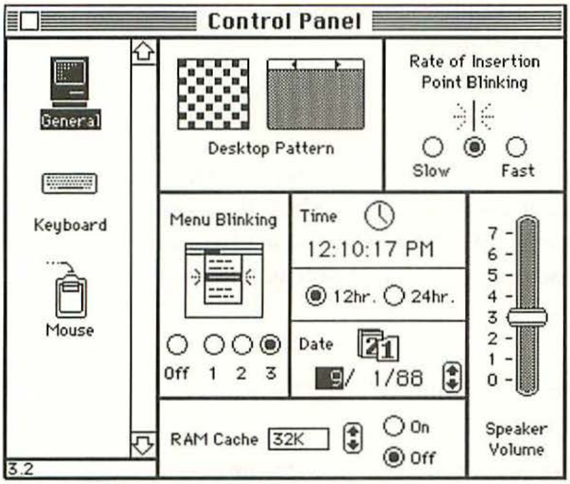

In this section we've looked at standard accessories supplied by Apple with the Macintosh. Many others are available from other software sources. (See Chapter 14.)

EDITING TEXT

For most Macintosh application programs, you will use the same general methods for entering and editing text. You can practice these operations using the Note Pad, since most system disks have one.

Choose Note Pad from the Apple menu.

If the first page has writing on it, turn to the next clean sheet by clicking the up-turned left corner.

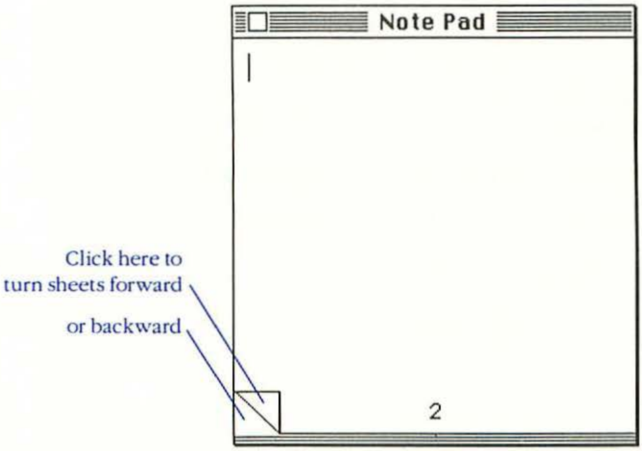

You should now see a blinking vertical bar at the top left of the Note Pad page; this bar marks the text insertion point. Anything you type appears at this point. This insertion point is used by all application programs for entering text.

Type two or three lines.

If you make a mistake, press the Backspace key and retype. Your text appears on the Note Pad, pushing the insertion point ahead of it. Notice how the words wrap around to the next line when you get to the edge of the page. You don't have to hit Return at the end of each line as you do on a typewriter; Macintosh software does it for you.

Inserting Text

Suppose you left out a phrase in the second line of the text you just typed, and you want to go back and insert the phrase where it belongs.

Move the mouse so that the pointer lies somewhere within the text.

Notice how the pointer changes shape as it moves onto the Note Pad, becoming an I-beam. You can use the I-beam pointer to mark where you want to insert the phrase you left out.

Move the I-beam pointer to anywhere on your second line of text and click.

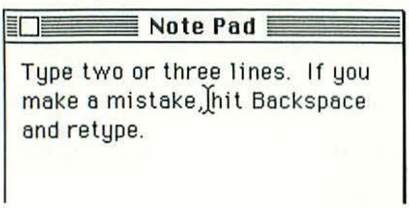

The blinking insertion point, which was at the end of your text, moves to the location of the I-beam pointer. If you now type in the missing phrase, it will be inserted in that location, pushing the insertion point and the existing text ahead of it. Practice moving the insertion point around your text.

Move the I-beam pointer to a blank area at the end of a line and click.

The insertion point appears at the end of the text in that line.

Move the I-beam pointer to the blank area following all the text you have typed and click.

The insertion point appears at the end of the last line.

Deleting Text

Click so that the insertion point appears somewhere within the text.

Press the Backspace key.

Using the Backspace key is the simplest way to delete something; backspacing removes the character to the left of the insertion point.

You can also use the Backspace key to delete more than one character at a time. But first, you have to use the I-beam pointer to select what you want to delete.

Move the I-beam pointer to the beginning of the text you want to delete.

Drag the I-beam pointer to the end of the text you want to delete and release the button.

Note that the text is now highlighted.

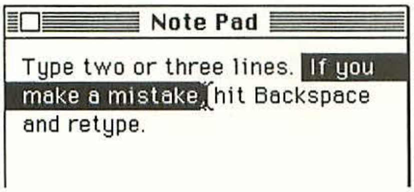

Press the Backspace key.

Gone!

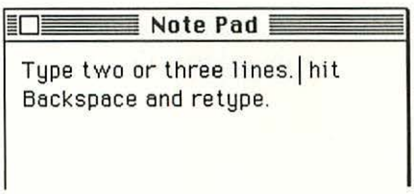

Once you've selected text with the I-beam pointer, you also have three other ways of deleting it:

Choose Cut from the Edit menu; *or*

Choose Clear from the Edit menu; *or*

Hold down the Command key (⌘) while typing *x*.

You can press the Command key and a letter instead of using the mouse to pull down and choose from a menu. The Command key equivalents for menu items are shown next to each item. From now on, this kind of instruction will be given simply as: Type Command-x.

Changing Your Mind

If you delete the wrong thing, the Macintosh gives you two ways to repair your mistake:

Choose Undo from the Edit menu; *or*

Type Command-z.

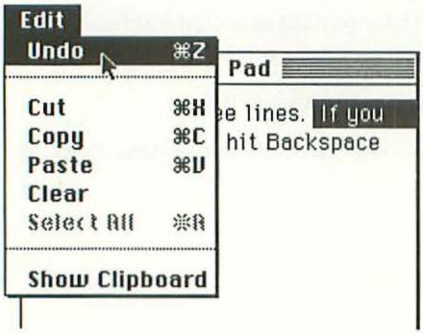

Either operation restores your last deletion.

Choose Undo again.

The selection is deleted again. Undo reverses whatever your last move was; if your last move was to restore a deletion, then Undo deletes the restoration. You cannot undo more than one previous step.

In some programs, the Undo menu item changes to reflect your last action. It might say Restore Row if you've just deleted a row, but the command will always be in the same menu location.

Erasing the Note Pad

Select the entire text by dragging the mouse down over all the lines.

Choose Cut or Clear from the Edit menu.

These actions erase the entire practice sheet.

Closing the Note Pad

Click the Note Pad's close box.

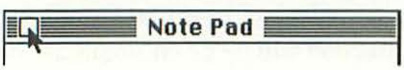

If you don't erase what you've written, your text is stored on the disk in a file named Note Pad when the window closes. The next time you use the Note Pad, you'll find your notes on the page where you typed them. Try it. To end this practice session:

Close the Note Pad window.

Close any other windows you might have opened.

Close the system disk window.

You can close windows one at a time by clicking each close box or all together by choosing Close All from the File menu.

Choose Eject from the File menu *or* Shut Down from the Special menu.

WORKING, MACINTOSH-STYLE

As you learn more about the Macintosh, you will find that the best programs follow a consistent procedure. Such programs will have no unpleasant surprises.

For most Macintosh operations, you always select something first, then choose what you want the computer to do with it. For example, you select a file, and then choose Get Info; you select a portion of text, and then choose Cut or Copy.

Not every function of every program works exactly this way. In the MacPaint program, for example, you select a tool first, and then you do something with it. Selections of an area within a MacPaint illustration follow the normal rule, however: Select first, then choose what to do with it.

Similarly, most application programs will also use the same two methods to get your attention: conversational dialog boxes and more urgent alert boxes.

Dialog Boxes

As you work with the Macintosh, application programs will at times put dialog boxes on the screen in response to your actions. For example, a dialog box might ask you the size of paper in your printer. Some dialog boxes merely ask a yes or no question.

In most dialog boxes, you respond by clicking buttons or items on a list. There are two kinds of buttons: round radio buttons, which indicate that only one choice is possible, or square buttons, which indicate multiple options. You can, for example, specify either automatic or manual paper handling for printing, but not both. But you can apply italics or boldface alone or both at once.

What happens next depends on the dialog box. Often, you click a button labeled *Yes* or *OK* when you are satisfied with your choices, or you click a button labeled *Cancel* to close the box and return to the previous step. You cannot proceed past a dialog box unless you click an appropriate button.

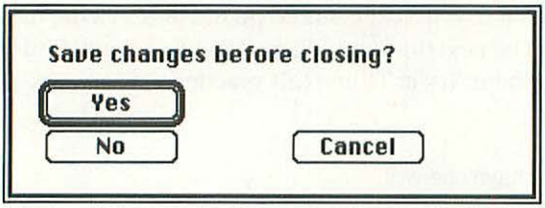

One special form of a dialog box is a minifinder—a listing of the files on the disk that were created by the application program you are using. A related form is a minilist, which might, for example, contain the font choices available for a particular program.

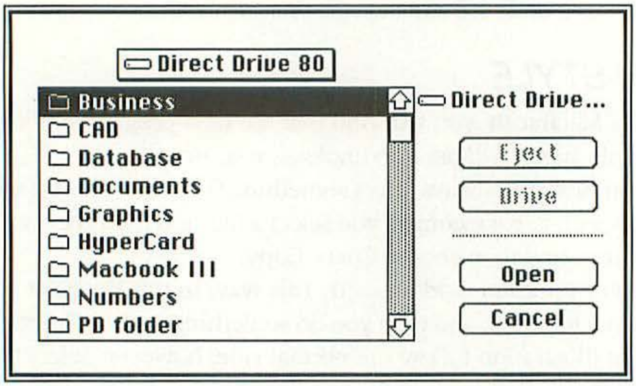

Alert Boxes

An alert box is similar to a dialog box, but it warns you of a potentially serious problem. For example, an alert box might tell you that your word-processing document is getting too long to store on the disk. In most cases, you must acknowledge the alert box by clicking an OK button.

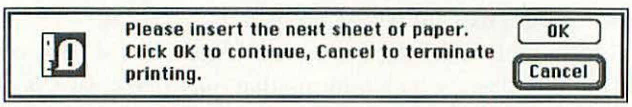

SO WHAT CAN YOU DO WITH A MAC?

Without the proper software, not much, for without software, a computer is hardly more than an overproduced doorstop. What kind of software you invest in depends, of course, on what you want your Macintosh to do. Today, hundreds of programs are available for the Mac; some of them may be just what you need to become more productive.

Word processors have made writing and editing anything from a short memo to a book far more flexible than with paper and pencil. Surgical revision can be done cleanly and precisely from a keyboard, rather than with scissors and tape; some word processors or page-makeup programs can create professional-looking formatting.

In the four years since the Macintosh was introduced, Macintosh graphics programs have progressed from relatively simple drawing tools to programs that can create complex engineering drawings, three-dimensional models, and color artwork.

At the same time, spreadsheet and graphing programs have become increasingly adept at doing business calculations and projections. Easy-to-use database programs for storing and sorting all kinds of data—numerical or otherwise—are becoming common, and programs for doing complex statistics are also starting to appear for the Mac.

As the Macintosh—and personal computing—evolves, business, artistic, and scientific application software will multiply and improve, offering innumerable choices where before there were none. For a survey of some specific software packages, see Section Three, "Working with Macintosh."

Sectio

Understanding the Macintosh

Now that more than one Macintosh model is on the market, how do you decide which one is right for you or which software will run on the model you have? By examining the nuts and bolts of Macintosh hardware and important accessories, this section provides some of the background you'll need to make these decisions.

3: How the Macintosh Works

Computers process information—numbers, words, graphs, pictures, or nearly anything you would call information. To be practical, a computer must take information in (input), manipulate the information (processing), get the information out (output), and save the information (storage).

This chapter discusses the basic functions of key Macintosh hardware components and the software that tells them what to do. Once you have a general idea of how these elements work together, you can find more detail about specific hardware and software in the following chapters.

BUILDING UP A MACINTOSH

Let's start with the most basic components: a keyboard to enter information, a microprocessor to manipulate it, and a video screen to display the output.

Suppose you want to type the letter *A*. When you press the *A* key, the keyboard generates an electrical signal corresponding to the letter *A*. This signal is sent to the microprocessor, which, with associated components, turns it into a different electrical signal and sends it on to the screen. These steps are physically accomplished by Mac's hardware: the keyboard, microprocessor, video screen, and connecting wires.

Computer programs, or software, control all the hardware. Software is nothing more than a set of instructions for the microprocessor; it enables the microprocessor to understand the keyboard's signal for the letter *A* and to create the pattern of dots that makes an *A* appear on your screen.

Memory

The hardware and software discussed so far constitute nothing more than a video typewriter—interesting perhaps, but not especially useful. A computer needs to be able to move, copy, and otherwise work with keyboard entries. To do this, it requires a kind of scratch pad—a place to keep keyboard characters while manipulating them.

Random-access memory (RAM)

The computer's scratch pad is an electronic storage area called random-access memory, or RAM, which resides on thin wafers of silicon called memory chips. RAM is fast; characters can be stored or retrieved in a microsecond (a millionth of a second). *Random access* means that the microprocessor can go instantly to any spot in the storage area for information, without having to start at any particular place, and can then jump forward or backward to another spot without having to read any information in between. A special area, called video RAM, is set aside as a map of the screen; every dot on the screen is mapped onto a bit in memory. Software controls this area to produce the images you see. (The Mac II's video RAM is separate from the main RAM.)

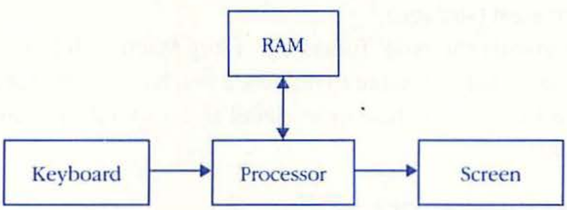

RAM has one major limitation; information stored in it is transient. When the power goes off, the information disappears.

All Macintoshes also have a tiny RAM apart from the main portion, called parameter RAM. This small section remains alive all the time, sustained by battery power. Parameter RAM stores the time, date, Control Panel preferences, and the way you've set up the modem and printer ports.

On the Mac Plus and earlier models, the battery, which generally lasts several years, can be removed through a small door on the back of the computer. Removing the battery for more than a few minutes erases parameter RAM; you will need to reset the clock and calendar afterwards. On the Mac SE and Mac II, the battery is soldered into the main circuit board; the computer must be disassembled to change it.

Read-only memory (ROM)

Programs must be in electronic memory to instruct the microprocessor, but not all software operates from random-access memory. Because RAM offers only temporary storage, some instructions—such as the program that tells the microprocessor what to do when the power comes on—can't stay there. These programs are held in another form of electronic memory called read-only memory, or ROM. ROM information is permanently stored on one or more silicon chips.

In all computers, ROM contains the initial instructions for starting the computer. The Macintosh's ROM also contains essential programs for controlling how a disk drive works, for interpreting input from the keyboard and mouse, and for drawing graphics or text on the screen.

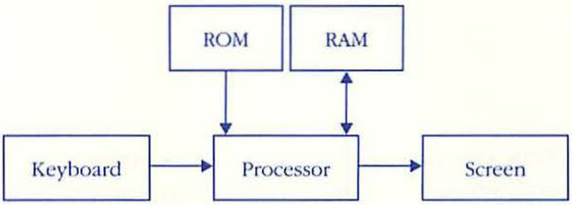

From the microprocessor's standpoint, ROM is simply another information source, like RAM, and even a little faster. But whereas RAM is transient, ROM is fixed. Once the computer leaves the factory, its ROM is permanent whether the power is on or off; the only way to change it is to replace the chip itself. Because ROM chips can be replaced, however, ROM is not as permanent as hardware, so ROM programs are sometimes called firmware.

Storage

To store large amounts of text, graphics, or other data for long periods, microcomputers use disk drives that can read and write information on magnetic disks, just as a tape recorder can play and record sound on tape. Like a tape recorder, a disk drive stores data as a series of magnetic pulses, except that on a disk the pulses are arranged in concentric circles. Disk drives come in two basic forms: Floppy disk drives store information on flexible magnetic disks; hard disk drives store information on rigid metal or glass disks.

Although a disk drive can read and write much faster than a human, it is far slower than RAM; finding a character on disk can take a disk drive from a tenth of a second (on a hard disk) to a second (on a floppy disk), compared with RAM's microseconds.

Because of this tremendous speed difference, all microcomputers normally use RAM for active work and disks for permanent storage, exchanging information between the two as needed. When you finish working with one chunk of information, or file, you tell the computer to store it on disk so that RAM is free to work with another chunk of information.

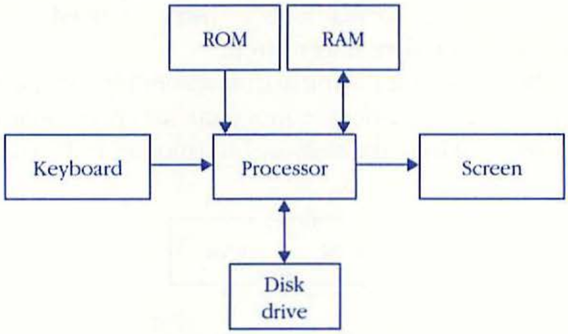

Like text or graphics, computer programs can be stored in disk files. For fast, effective operation, a program must be read (temporarily copied) into RAM before use. When you insert a MacWrite disk, for example, and start the program, the computer transfers a copy of MacWrite instructions from the disk into RAM. If you quit MacWrite and change to MacPaint, the MacWrite instructions in RAM are replaced with MacPaint instructions. In both cases, the programs remain stored on the disk as well.

BITS AND BYTES

Computers can only process information they understand, and they only understand electronic signals that are either on or off. A computer processes information in the form of individual on/off signals, or bits, coding each bit as 1 or 0.

One bit can't convey much information, so a computer strings many bits together. A single character (a letter of the alphabet, a number, or a punctuation mark) is coded by eight bits in sequence, or one byte. The letter *A*, for example, is coded as 01000001, *B* is 01000010, and so on. Each hardware component—keyboard, memory, disk drive—codes the letter *A* according to this standard, called ASCII (American Standard Code for Information Interchange). Nearly all microcomputers understand ASCII.

The information content of a single byte is still limited, so information is usually measured in kilobytes. A kilobyte, or 1 KB, is 1024 bytes. Although *kilo* ordinarily means 1000, a kilobyte is not an even 1000 bytes because the computer's counting system is based on the number two, not the number ten. The power of 2 nearest 1000 is 2 multiplied by itself 10 times (2^{10}), or 1024.

Disk file size is customarily measured in kilobytes. A file 6 KB long contains about four pages of text, or 6144 characters (a typical double-spaced typewritten page holds about 1500 characters). A 6 KB file doesn't have to consist of characters, though; it can be a program or a picture of equivalent size.

A few other measures: 1024 kilobytes equal 1 megabyte; 1024 megabytes equal 1 gigabyte (about one billion bytes); 1024 gigabytes equal one terabyte (about one trillion bytes). For simplicity, it is often said that a megabyte equals 1000 kilobytes, a gigabyte 1000 megabytes.

THE BUS

Coded information travels through the computer on a bus, a set of wires that serves as a data highway linking the computer's components together. Each component communicates with the microprocessor via the bus.

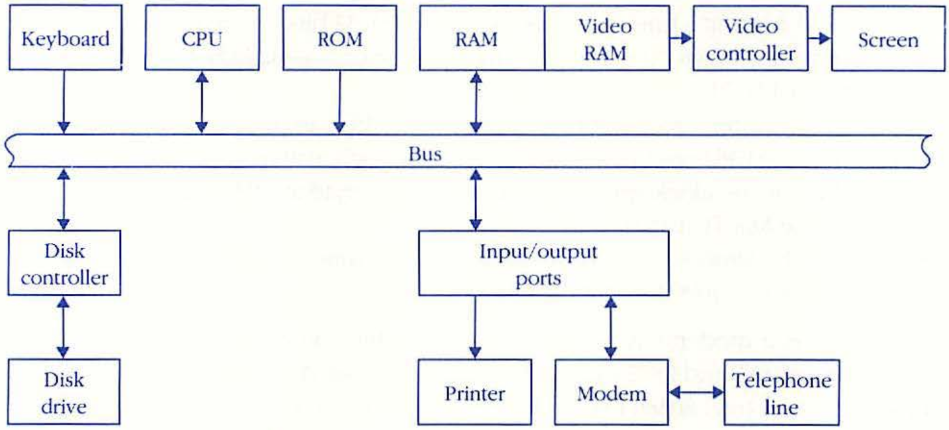

The bus carries three main kinds of information: One group of wires conveys data, such as the coded letter *A;* another group carries the address of the component to which the data are headed. Each component accepts only information addressed to it; for example, information intended for the printer will not inadvertently be accepted by the disk drive. The third kind of information consists of timing signals, which synchronize everything connected to the bus to send signals and listen at the correct time.

On Macintoshes before the SE, the bus was buried inside the computer, and accessory circuits could be connected to it only with difficulty. On the SE, a single connection called an expansion slot has been added to the bus so that you can quickly and easily attach an accessory circuit board (also called a card). The Macintosh II has six of these expansion slots; one is normally occupied by the video circuit board, leaving five free for other accessories. (For more expansion slot information, see Chapter 7.)

Most of the action on the bus is orchestrated by the computer's central microprocessor.

THE CENTRAL PROCESSING UNIT (CPU)

The heart of every microcomputer is a single integrated circuit chip—the microprocessor, or central processing unit (CPU). The CPU chip in the small Macs is a Motorola 68000; in the Mac II, it is a Motorola 68020.

By itself, the CPU is powerless. To do anything useful, it carries out, one step at a time, the instructions provided by software. An instruction might read: "Take the information stored in memory location 8, add 1, and put it in location 20." A more complex set of instructions might be: "Take the character the keyboard has placed on the bus and put it in memory location 300." Each step is simple; computers do useful work because they can perform millions of steps in rapid succession.

A CPU's power depends on four factors:

- How much information the CPU can work on at once, measured two ways:
  - —How many bits (called input/output, or I/O, bits) it can take from and put onto the bus at a time—16 bits for the 68000, 32 bits for the 68020.
  - —How many bits it processes internally at one time—32 in both the 68000 and the 68020.
- How many different kinds of instructions the CPU can perform.
- How fast the CPU operates—how much time each instruction takes, usually specified by the clock speed; the small Macs operate at 7.83 megahertz (MHz), the Mac II at 15.67 MHz.
- How much memory the CPU can address at one time—16 MB for the 68000, four gigabytes (4096 MB) for the 68020.

The 68000 is a moderately powerful chip, roughly comparable to the Intel 80286 used in IBM's PC/AT and PS/2 models 50 and 60. The 68020 is more powerful, roughly comparable to the Intel 80386 used in the IBM PS/2 model 80.

Memory-addressing capacity used to be a severe limitation in early microcomputer CPUs. Many could not address more than 64 KB at a time; some could address 1 MB. Now that both the 68020 and 80386 can address 4096 MB, a microcomputer's overall power is affected less by CPU memory-addressing capacity than by how much memory can physically fit inside the computer and the memory-addressing limitations imposed by particular software.

THE MANY FORMS OF SOFTWARE

Some of the software that controls the Macintosh is built in, and some of it is available on disks. Some essential software tells the microprocessor how to read a disk; other software can draw a picture of a disk. You interact directly with some software, while other software makes that interaction possible.

ROM Software: Key to the Visual Interface

Apple's Lisa and later the Macintosh were the first microcomputers to have a comprehensive collection of programs in ROM. It is these programs that provide the power behind Apple's visual interface.

ROM software controls the interface, drawing key parts of what you see on the screen, monitoring the mouse, and so on. It therefore defines the way that you deal with application programs—the word processors or spreadsheets you do real work with. ROM software makes up a kind of programmer's tool kit, to be used by professional and amateur programmers when designing other Mac software. This ROM-based interface is the reason that different Macintosh applications work much the same way.

Dozens of programs exist in ROM; the following are among the most important:

- QuickDraw draws complex graphics on the screen quickly.
- The Font Manager uses QuickDraw to draw characters on the screen.
- The Event Manager keeps track of what you do with the mouse and keyboard.
- TextEdit is a collection of text-editing routines.
- The Window Manager draws and controls windows on the screen.
- The Control Manager creates and monitors dialog boxes and which buttons you choose within the boxes.
- The Menu Manager creates and monitors the pull-down menus.
- The File System creates and controls files in memory and on disk.
- The Resource Manager keeps track of resources needed by programs, loading them into and removing them from RAM as needed.

Apple has developed several versions of ROMs for the different Macintosh models. All contain the programs sketched above, but later Macs have ROMs with many more programs. The ROMs are:

- The original 64 KB ROM used in the original 128 KB Mac and the Mac 512K.
- The 128 KB ROM used in the Mac Plus and Mac 512KE, which improved operational speed and display quality.
- The 256 KB ROM used in the Mac SE, which operates faster still.
- The 256 KB ROM used in the Mac II, which, in addition to the features of the previous ROMs, supports color displays and other accessory hardware.

Although most 1987 software can run with the original 64 KB ROM, more and more programs require routines that were first built into the 128 KB ROM, such as support for characters of fractional pixel width (5½ pixels, for example). Some programs appearing in 1988 and later will be able to work only with the Mac II ROM. As Macintosh software proliferates, it will be increasingly important to check which Mac model a program can run on; some will run on many models but will offer all their features only on the latest Macintoshes.

The Operating System: A Traffic Cop

A fundamental program called the operating system acts as traffic cop, monitoring and directing all Macintosh operations. It manages everything in memory and keeps track of information going to and from each component, whether the disk drive, printer, keyboard, or screen.

In conventional microcomputers, the operating system is read in from a disk and stays mostly in RAM when the computer is working. Some popular operating systems are Apple DOS (Disk Operating System) for the Apple II; MS-DOS (Microsoft Disk Operating System) or PC-DOS for IBM PCs and compatibles; and OS/2 for the newest IBM microcomputers.

In contrast, most of the Macintosh operating system (which has no name) resides in ROM. The rest is stored in a file named System on the disk you use to start the Mac (any disk with a System file on it is called a system disk). Portions of this file are read from the disk into RAM when you first turn on the computer; they add to or modify the instructions in ROM or both.

The System file

The System file contains many resources.

Utility programs perform specific tasks that support the overall operation of the computer. When you select a file icon and then choose Get Info from the File menu, for example, you are actually starting a small utility program that checks that file and displays information about it. Some utilities can be nearly as important as the operating system; others are simply handy to have available.

Desk accessories, chosen from the menu under the Apple symbol at the far left of the menu bar, include conveniences such as a clock, a calculator, or a note pad. Because these programs are usually small (less than 32 KB), they can share RAM space with an application program.

Font information dictates the design, size, and style of the letters you see on the screen. Several fonts are essential to the Mac's operation—the ones you see on the Desktop and in the main menus, for example. Many additional type styles and sizes are also available from this file when you are using word-processing, graphics, and other programs. To create these, the operating system transfers font information from the system file to RAM. If you change type styles or greatly change type size, you may have a short wait while the Font Manager program (in ROM) goes back to the disk to bring the new information into RAM.

Because information about each font takes up considerable disk space, you may want to use the Font Mover program to store rarely used fonts on a separate disk. (Desk accessories and fonts can also be stored with specific application programs rather than in the System file. For more about fonts see Chapter 12.)

Messages residing in the System file consist of both warnings and advisories. For foreign versions of the Macintosh, Apple changes this resource to give messages in the appropriate language. (The ROM programs themselves contain no text in any language.)

The Finder: Keeping Track of Disk Files

Another important program stored on disk and read into RAM when you first turn on the Mac is the Finder. The Macintosh Finder includes many functions traditionally performed by a computer's operating system. It handles most operations that involve disks: creating the disk window with its file icons, copying files, copying disks, and so on. The Finder doesn't work alone; it uses many programs in ROM for actual disk access, in effect acting as liaison between you and the ROM programs that control the disk drives.

Each disk has a directory that acts as its table of contents. The directory contains file headers, which hold information about each file on that disk. When you insert a disk, the Finder puts this directory information into RAM, where it remains even if you change disks. You can display on screen the directory of an ejected disk, but if you actually want to work with a file from it, the Finder asks you to insert the disk.

Apple revises the Mac operating system and Finder from time to time, usually to add new features, but also to correct errors ("bugs"); the updates are usually available from Macintosh dealers.

The DeskTop file

For each disk, the Finder creates a hidden DeskTop file to hold information about each file. The DeskTop file "remembers" such on-screen details as the icon associated with a file and the size and format of the disk's window on screen. It also notes whether a file is an application program or a document; if it's a document, the DeskTop file checks the file's header for the application program that created it and keeps track of the icon images required by the program. (See Chapter 14 for more details on file structure.)

Clock/Calendar

Every Mac has a battery-operated clock/calendar whose time and date are read into RAM. Every time you create or modify a disk file, the date is automatically stored with the file header. This time-and-date keeper drives all software that displays current time.

Application Programs: Your Primary Tools

You use application programs for doing work—a word processor to produce written documents, a paint program to create illustrations, and a spreadsheet program for financial calculations. Permanently stored on disk, application programs are transferred by the disk drives (loaded) into RAM when you need them. The latest version of the Macintosh System file and Finder lets you read several applications into RAM at a time so that you can switch among programs without waiting for the disk drives. How many programs you can read in depends on how much RAM your Mac contains and how big the applications are.

Some applications need not be completely loaded into memory before you can use them. When you start such a program, you read in a core program that stays in RAM while other program components, called overlays or transient code, remain on the disk until needed. As you choose commands—a sorting routine or trigonometric calculation, perhaps—the core program brings needed overlays into RAM; each new portion replaces other overlays not then in use.

Although the overlay procedure lets you use powerful programs that can otherwise exceed RAM capacity, it also slows operation. But as Macintoshes gain more memory, overlays are becoming less necessary. A complex application contained entirely in RAM runs much faster than with overlays.

THE MANY USES OF RAM

A Macintosh has far more activity going on in RAM than a conventional microcomputer does. RAM holds a lot of information:

- Video memory
- Parts of the operating system
- Utility programs
- Desk accessories
- Current font data
- Icon images
- The Finder
- Disk directories
- Application programs

RAM also contains two forms of your data:

- Data used in your applications
- Data on the Clipboard

Data: Your Information

You can enter data into the Macintosh through the keyboard, mouse, or disk drive, or, from another computer, over a telephone line or through a network linking the computers together.

From an application program's standpoint, the data source doesn't matter because, in most cases, the program will put the data into RAM before beginning work. Some programs must have all the data in RAM at one time; others can read some data into RAM and leave the rest on disk, swapping chunks as needed.

The Clipboard

Once some or all of your data are in RAM, whenever you cut or copy anything from a document or file, that information goes into the Clipboard, an area of RAM set aside to act as a holding area for information exchange between programs. For example, you may want to cut a series of numbers from a Multiplan spreadsheet and paste it into a MacWrite document so that you can include a financial statement in a memo.

You can store text, a drawing, or numbers in the Clipboard, but you can store only one item at a time. If you need to store more items, paste the Clipboard contents into the Scrapbook; this frees the Clipboard for another item.

When you quit a program, the contents of the Clipboard are written on the disk in a file called Clipboard.

WHAT HAPPENS WHEN YOU START A PROGRAM

To tie together the hardware and software components discussed in this chapter, here is a brief outline of what happens when you start MacWrite. The outline is not complete, and events don't occur quite so linearly, but it will give you some idea of how the Macintosh system works.

When you turn on the Macintosh, a ROM program called Boot tells the CPU to check whether a system disk has been inserted into the floppy disk drive and, if not, whether a hard disk drive with a System file is running. (The term *boot* comes from the idea that the computer is pulling itself up, or on, by its own bootstraps.) If no disk is ready, the Boot program puts an image of a disk with a question mark on the screen.

Once a System file is found, the Boot program instructs the disk-controller circuitry to send electronic signals to the disk drive to move the disk drive head to the disk's outer edge and begin transferring information from the disk into RAM. First, the part of the System file containing the RAM portion of the operating system is read from the disk. The Mac then reads in the Finder from disk and creates the Finder display, the Desktop.

To find out what program and data files are stored on the boot disk, you select the disk image by clicking it with the mouse; you then choose Open from the File menu (or double-click on the disk icon). The Finder creates a window showing the file icons and names. You move the mouse so that the pointer is, say, over the MacWrite icon. The Event Manager (in ROM) detects the mouse's position. You double-click on the icon. The Event Manager tells the Finder about the clicks. The Finder checks the pointer location and concludes that you want to start the MacWrite application program.

The Finder checks the disk directory for MacWrite's location and passes the location to the disk controller, which starts turning the disk and moves the disk drive head to the beginning of the program file.

As the disk drive head reads this file, the disk controller puts the information on the bus. From the bus, the information passes into RAM, to space allocated by a ROM program called the Heap Manager.

Once in memory, MacWrite begins changing the screen. It replaces the Finder menu bar with the MacWrite menu bar. Almost simultaneously, the Window Manager (in ROM) puts a window on the screen, complete with scroll bars and title.

The Font Manager (in ROM), which has been busy displaying the text on the screen, also checks over the System file to see which fonts it contains. It passes on the number and available sizes of each font to the Menu Manager (in ROM), which in turn sets up the Fonts menu.

As you can see from this incomplete outline, even something as simple as starting a program requires an enormous number of steps. That the procedure works at all is amazing; that it works so well is a tribute to the thousands of engineers and programmers who, during the last 50 years, have made computers possible.

4: The Video Screen

The Macintosh display screen is a cathode-ray tube (CRT), like an ordinary television screen but with much higher resolution and a much sharper image. The image is made up of tiny dots called picture elements, or pixels. On the Mac II, the separate CRT unit is called a monitor.

Almost every computer display screen is rectangular. Those that are wider than they are tall are called landscape screens; those taller than they are wide are portrait screens. Eventually, screens that can be rotated 90 degrees—for either portrait or landscape orientation—will appear; software will adjust the image as required.

MAPPING THE SCREEN

All the information needed to generate a screen image is stored in a type of random-access memory called video RAM; this can be a portion of the main RAM (as in the small Macs), or it can be located on a separate circuit board called a video card (the Mac II). The size of video RAM depends on screen characteristics. In the small Macs, it takes up a little more than 21 KB (175,104 bits per screen divided by 8 bits per byte divided by 1024 bytes per kilobyte). More RAM is needed by larger screens with more pixels and by screens that display shades of gray or color, because, on top of all the other image information, these screens must also store gray or color data for each pixel.

Each pixel on the screen corresponds, or is "mapped," to a bit in video RAM; in computer jargon, the Macintosh screen is bit-mapped. To generate the screen image, Macintosh system software takes the display information created by an application program, computes which pixels should be lit up, and puts all this data into video RAM.

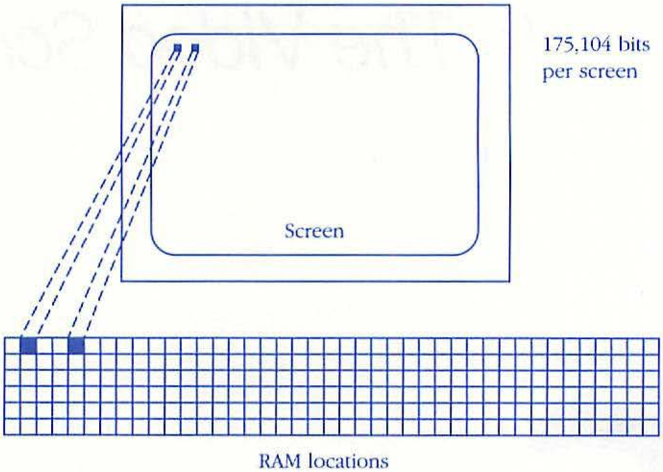

Bit-mapped screen: Each dot on the screen corresponds to a bit in memory.

Drawing bit-mapped screen images calls for a lot of computing. Every time you change something on the screen, the computer must recompute each pixel of the changed object or area on the screen; unless a program is poorly designed, the computer usually does not recompute stationary items.

CREATING THE SCREEN IMAGE

Once all the screen bits have been computed and stored in video RAM, a video-controller circuit takes this information and creates a video signal suitable for the screen. The video signal controls a beam of electrons generated by an electron gun, turning it off and on for black and white or modulating it for grays. (Historically, electron beams were called "cathode rays," after the device, a cathode, that produced them.) The electron beam, pointed toward the inner side of the screen, passes between electromagnets that direct the path of the electrons. When the beam is on, the electrons strike a phosphor coating on the CRT face, and light is produced for that particular pixel.

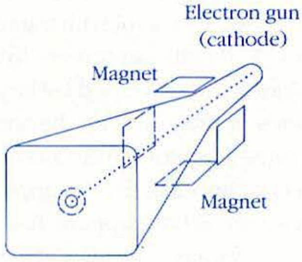

Electrons strike phosphor, creating light.

On the small Macintoshes, the video controller is built into the main computer circuit board. On the Macintosh II and all large-screen accessories for any Mac model, the video controller comes on the video card, which contains video RAM and can contain its own processor and ROM as well.

On the small Macs, the images you see are made of pixels that are either on or off. These Macintoshes, and the external accessory display screens made for them, cannot show true grays, although objects can be made to look gray by turning off alternate pixels. On a Mac II with a monochrome display (one color or black-and-white), each pixel can show true grays—up to 16 levels at a time with Apple's Macintosh II Video Card (4 bits stored per pixel in video RAM); with the additional video memory on the Macintosh II Video Card Expansion Kit, up to 256 gray levels (8 bits stored per pixel) can be displayed simultaneously. Although 256 shades of gray may be more than a CRT can display distinctly, having so many gray steps produces smoothly graded tones when your work requires them. (The Mac II also supports color; see page 60.)

To keep an image constant, the beam of electrons within the CRT sweeps across the screen repeatedly in a series of successive scan lines. Macintosh video circuitry refreshes the screen image more than 60 times each second, often enough that your eye and brain perceive a continuous image.

The Macintosh has a nonstandard video signal; it displays lines 50 to 100 percent faster than standard video systems, and unlike ordinary televisions, which generate images by drawing every other line 60 times a second, the Macintosh draws *every* line 60 or more times a second (in jargon, the Macintosh produces a noninterlaced image). These differences mean that you cannot use ordinary television equipment to directly record or broadcast the Macintosh video signal—to make a videotape for training, for example. Several companies, however, offer converters that transform the Macintosh's video signal into a form compatible with standard video equipment; see Chapter 22 for more information.

BIG SCREENS

The original 9-inch Macintosh screen has a 4¾-by-7-inch (12-by-18-centimeter) image area—small by most standards. (The exact size depends on the settings of horizontal and vertical size controls in a particular display.) Since 1985, however, several manufacturers have come out with larger screens: Some simply magnify the image without an increase in the number of pixels, and some actually have more pixels.

Large screens with the same number of pixels as the original Mac screen are useful for those whose visual acuity, focusing ability, or both are restricted or for a group of people who must look at a single display. Such screens consist simply of a larger CRT or a projection video system; connecting them to a small Mac requires opening the case and tapping the video signal. (See Chapter 22.)

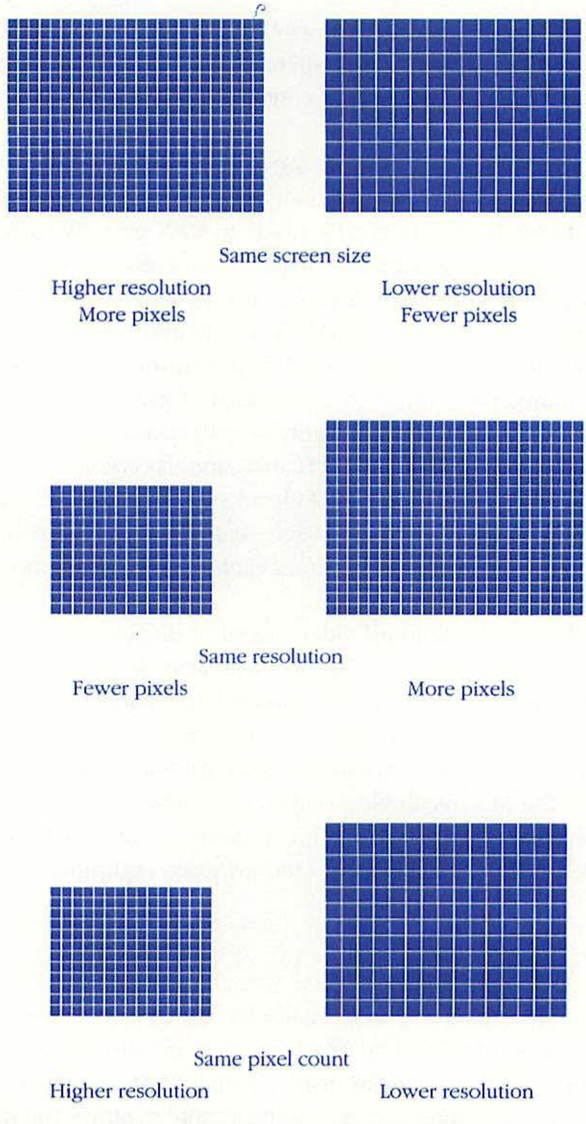

The relationship of screen size, resolution, and pixel count. Each dot represents a pixel.

For graphic arts or desktop publishing, big screens with more pixels are far more appealing. Screens that display one to four full pages are now available for the Mac Plus, Mac SE, and Mac II. But although these screens have a higher pixel count, they do not necessarily offer higher resolution than the original Mac screen. Often misused, the term *resolution* refers to the number of dots, or pixels, per inch (dpi) on the screen; pixel count is simply the total number of pixels. Two screens of different sizes but with the same number of pixels will therefore display images at two different resolutions, and two screens of different sizes but the same resolution will have different numbers of pixels. (See Table 4-1.)

TABLE 4-1. SOME REPRESENTATIVE MACINTOSH II VIDEO SCREENS[1]

Screen	*Pixels horiz.*	*Pixels vert.*	*Pixel count*	*Resolution (dpi)*	*Gray levels*	*Colors*	*Screen size[2]*	*Page size[3]*	*Orientation*
Original Macintosh screen	512	342	175,104	72	2	None	9	0.35	Landscape
Apple Video Card for Mac II	640	480	307,200	80	16	16	12	0.63	Landscape
with expansion kit	640	480	307,200	80	256	256	12	0.63	Landscape
E-Machines Big Picture	1024	808	827,392	82	256	None	17	1.68	Landscape
MicroGraphic Images	1024	900	921,600	75	2	None	19	1.88	Landscape
Nutmeg Information Systems	720	900	648,000	90	2	None	15	1.32	Portrait
Radius Full Page Display	640	864	552,960	77	2	None	15	1.13	Portrait
Sigma LaserView[4]	1664	1200	1,996,800	150/115	4	None	15/19	4.06	Landscape
at lower resolution	832	600	499,200	75/65	4	None	15/19	1.02	Landscape
SuperMac Spectrum	1365	1024	1,397,760	72/85	256	256	16/19	2.84	Landscape

[1] All specifications are subject to change.
[2] Diagonal, in inches.
[3] Number of pages displayed. A page is defined as 600 pixels wide by 819 pixels high, for a total of 491,400 pixels. For most software that drives the LaserWriter printers and allows for borders, this corresponds to a letter-size page.
[4] Screen can be switched between two resolutions.

Connecting Big Screens

Most large screens can be easily connected to the Mac SE and Mac II, and some will work with the Mac Plus as well. Even though large screens were first developed for the Mac Plus, installing one on it is not easy. If your work needs a big display, you should get a Mac SE or II. In most cases, the video monitor is the same for all Macintosh models; only the video-controller card differs. Thus, if you have a big screen connected to an SE and you upgrade to a II, you can preserve most of your investment simply by getting the proper controller card.

Connected to a Mac Plus or SE, some big screens displace the original screen altogether, giving you a single large display. Others, in contrast, run at the same time as the built-in Macintosh display. When placed next to each other, the two screens form a disjunct display shaped like an inverted L, with a large work area on the big screen and room for desk accessories or an extra window on the original Mac. Although clever, this two-part arrangement isn't especially practical; the disjointed images make the Mac screen ineffective as a real extension of the large screen, and many people simply ignore it.

The Macintosh II takes the concept of multiple displays even further, supporting several screens, each connected to its own video card, all running at once. For certain tasks, multiple displays can work well—if you need to compare several large drawings, for example. Or you can combine a monochrome screen for text with a color screen for graphics. For extra-large images, multiple screens on a Mac II can also be functionally joined together, but a single bigger screen is easier to use than several smaller screens together.

Working with Big Screens

Particularly if you are already accustomed to a small Macintosh screen, living with a big screen requires some adjustments. In the first place, large displays are physically bulky and heavy. Less obvious is an operational problem—the sheer distance between the bottom of the screen and the pull-down menus at the top. If you are working at the bottom of the screen, it can take a while to move to the top, choose a command, and move back to your work area.

The software for the mouse helps alleviate this problem by sensing quick mouse movements and interpreting them as a long-distance move on the screen. Programs that have keyboard commands for scrolling and other common functions also make working with big screens easier. In the long run, big screens will encourage development of "tear-off" menus that can be moved to wherever you are actively working. MacPaint, version 2.0, already has some of these features, and Radius (Sunnyvale, CA), a manufacturer of big screens, supplies utility software to create tear-off menus in some existing application programs. Other software tools let you reset the size of both the pointer and the system font, which is normally 12-point Chicago. These tools can replace the system font with a larger or smaller typeface, thus changing the size of the menu bar at the top of the screen.

Pixel count and resolution also affect working with a big screen. In terms of perceived image sharpness, the more pixels and the higher the resolution, the sharper the image. But higher resolution has some significant side effects. For their displays, most Macintosh programs assume a fixed pixel size based on the original 72-dpi screen. Rulers on the screen follow this scaling; one inch on screen corresponds to one inch on paper. But on a screen with 90 dpi, for example, the same ruler appears 20 percent smaller, as do menu items, the Apple logo, and any other object with a fixed number of pixels. Unless the software can make an adjustment, if you want a screen to show the same scale as a paper drawing, be sure it has 72 dpi.

Partly to address scaling issues, some manufacturers offer screens that operate at more than one resolution and pixel count. Sigma Designs' (Fremont, CA) LaserView screen lets you switch between two resolutions and pixel counts—from nearly 2 million pixels and 144 dpi to a half million pixels and 72 dpi, the same as a standard Macintosh screen.

An unexpected problem can arise if you use a dual-resolution screen like this or two screens with distinctly different resolutions. To position the mouse accurately, you rely on both kinesthetic feedback from your arm and wrist muscles and visual feedback from watching the pointer on the screen. If the visual feedback changes scale, you might find precise mouse movements a little awkward each time you change resolutions.

Dual-resolution screens can pose still another problem. The size of a single pixel created by the beam of electrons striking the screen is usually equal to the reciprocal of the resolution (on a 72-dpi screen, one pixel is $1/72$ inch), so the scan lines create a continuous image without gaps or overlap. If pixel size remains unchanged on a dual-resolution screen, you might see gaps between the lines at low resolution or overlap at high resolution.

If you buy a big screen for a Mac II, be sure that it and its supporting video card, like Apple's, can display at least 256 levels of gray. Many early video cards for big screens could display only black and white, and many new programs require grays.

IMAGE QUALITY

On big or small screens, perceived image quality is affected by several factors, including resolution. Although manufacturers quote specific resolutions for their screens, actual resolution can easily vary by 5 percent or more because of variations in and aging of components.

Screens rarely look absolutely linear; one side or the top or bottom is often somewhat distorted. Most monitors contain controls for adjusting linearity (like the horizontal and vertical height/size controls on a television set), but these controls are often internal and normally set only by a technician. Even with adjustment, a few percent nonlinearity is common, if not usually bothersome. If your requirements are exacting, check the specific monitor you plan to buy and get it adjusted as accurately as possible; readjust it annually. Alternatively, you could get a flat-panel screen, which is completely free of nonlinearities (see page 62).

The crispness of a screen image depends on the computer's video circuitry and the CRT. The larger the screen or the higher the resolution, the more difficult it is to produce a really crisp image with sharp, black pixels. In nearly all screens, the pixels in the center are the most clearly defined. All screen images appear a little soft at the corners—large-screen images more so than small ones.

Everything else being equal, a screen whose face is flat is more desirable than one whose face is curved. Truly flat CRT screens are just beginning to appear; Zenith (Glenview, IL) has come out with the first truly flat-surface color CRT. A flat image is less distorted than a curved one and can be seen from a wider angle. Flat-faced screens also have less glare because the surface catches fewer reflections.

COLOR SCREENS

Color monitors are much heavier and more expensive than monochrome monitors of the same size, but they work fundamentally the same way. Instead of one electron gun, most color CRTs have three; they also have three phosphor coatings on the CRT face, one for each primary color (red, blue, green). These are laid down in triads of either rectangular stripes or triangles of three dots. The beam of electrons from each gun strikes, and thereby illuminates, only one color phosphor. (All color monitors use three phosphors; in some designs a single electron gun is switched rapidly among all three colors.) Colors other than red, blue, or green are produced by additive mixing of the primary colors.

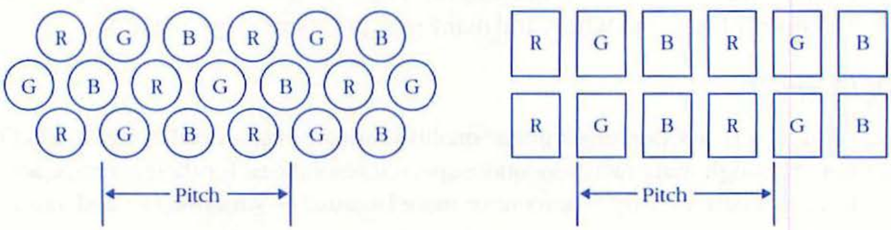

Triads of red (R), green (G), and blue (B) phosphor dots (left) or rectangular stripes (right). The pitch of a color screen is the distance between repeating triads.

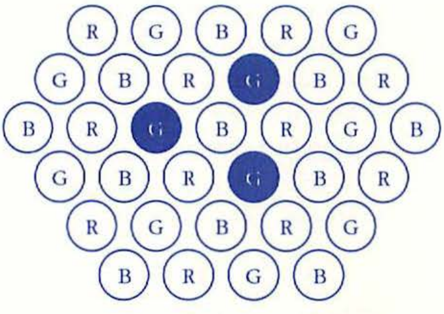

A pixel is usually made up of several phosphor dots. In this example, a green pixel consists of three phosphor dots.

Before the Mac II, there was no practical way to graft a color CRT onto a Macintosh. (A few companies attempted to build expensive, complex color monitors that could be driven as output devices, somewhat like a printer. But without special software, such displays cannot work interactively as the Macintosh normally does.) Orchid Technology (Fremont, CA) produces a color video board for the Mac SE. Because the SE's ROM does not fully support color (unless Apple releases an upgraded ROM), Orchid's board requires special software drivers for each program that you want to run in color.

The Mac II, on the other hand, was designed for color displays. If color is important to you, you will be best off with a Mac II, despite its cost. Apple sells a 640-by-480-pixel color monitor made by Sony, and other companies build even larger displays; the big Sony monitors sold by SuperMac (Mountain View, CA), for example, display 1024 by 768 pixels. Standard color television monitors will not work with a Mac II because they can only display 242 noninterlaced lines; they cannot handle the Mac II's 480-line noninterlaced image.

Apple's monitor and video card can show 16 to 256 colors simultaneously (either 4 or 8 bits per pixel) from a palette of more than 16 million. To get such a large color palette, the Mac II video cards send analog, or continuously variable, video signals to the color monitor; fewer colors are possible with digital, or discrete, signals. Even 256 colors, however, are not enough to create a truly realistic image. By scattering pixels of existing colors—a process known as dithering—some software can simulate colors on the screen that are not among the 256. Close up, the dithered area looks mottled; from a normal viewing distance, the color you see will be the average of the scattered pixels.

SuperMac and RasterOps (Cupertino, CA) both produce color video boards that store 24 bits of color information per pixel and can thus display more than 16 million colors simultaneously for a realistic image without dithering.

The Control Panel desk accessory on the Mac II allows you to choose how many colors (or shades of gray) you want to display. A huge number of colors on screen is not always an advantage, however, because the computer must compute more to display more colors. Having more colors slows scrolling and other functions that require recomputing the screen. A separate video processor (see page 64) can help diminish this performance penalty.

Compared with monochrome monitors, color displays suffer from a variety of difficulties. First, the three primary colors must exactly converge, or align, across the entire display to produce a clean image, and this is hard to achieve. To help out, the Mac II Control Panel includes a convergence test, which can be used along with special controls on the color monitor to adjust color alignment. The best, and most expensive, color monitors incorporate a special memory-correction circuit that divides the screen into 16, 64, or more zones and keeps track of convergence and linearity errors for each zone. As the electron beam sweeps across the screen during operation, the CRT circuitry corrects any errors in each zone. Such monitors are rare, but they are worth the expense if your requirements are exacting.

Second, color screens are inherently less sharp than monochrome displays because of the phosphor triads of primary colors. When everything is perfectly aligned, the size of a color triad becomes an important limiting factor for sharpness. The distance between color triads is called pitch; the smaller the pitch, the sharper the image. In monitors where the phosphors are arrayed in triangles, vertical pitch equals horizontal pitch; in those where the phosphors lie in vertical stripes, vertical pitch is set by the scan lines of the electron beam. Display manufacturers normally quote horizontal pitch.

Although the best color monitors perform credibly, they are still less sharp than a good monochrome monitor—and they are much more expensive. For a given investment, most people will probably find it more useful to get a larger or higher-resolution monochrome monitor than a color display.

FLAT-PANEL SCREENS

CRTs offer a bright, high-contrast image, but they are bulky, fragile, and power hungry. In contrast, flat-panel displays are shallow and take up little space; some consume considerably less power, and all are free of the nonlinearities and regional fuzziness that plague CRTs. Flat-panel screens are most common in transportable and portable computers, including the few portable Macintoshes on the market. (See Chapter 18.)

Three flat-screen technologies are now widespread. In all of them, a precisely laid-out array of electrodes switches an equally precise array of pixels on and off.

Liquid crystal displays (LCDs) are the most common flat-panel screens. The liquid crystal material consists of linear organic molecules that, at room temperature, exist in a phase intermediate between solid and liquid. Sandwiched between two flat sheets of plastic or glass, the molecules are arranged in a parallel array that polarizes light. The front flat sheet (closest to the viewer) is a fixed polarizer oriented along the same axis as the liquid crystal so that light passes through both the sheet and the crystal. When an electric field is applied to a liquid crystal, the molecules twist, changing their axis of polarization. Two kinds of LCDs are common today: nematic liquid crystals, which are twisted 90 degrees, and supertwist crystals, which are twisted 270 degrees. In either case, the light becomes cross-polarized with respect to the front sheet, so it no longer passes to the back reflective sheet. Where the molecules have been electrically activated, the viewer sees a dark spot. At any given instant, only one pixel is active; the controller circuitry activates each pixel in turn by scanning rapidly and repeatedly through the entire display.

When an LCD is inactive, its polarization axis is the same as that of the fixed polarizer in front of it. Light passes through both layers, reflecting off the mirror behind the LCD and returning to the viewer.

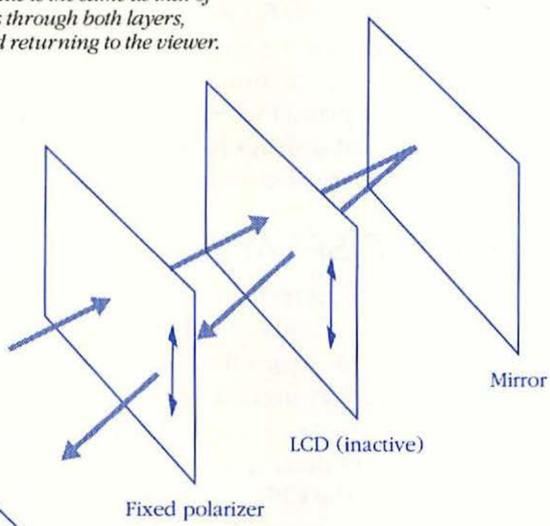

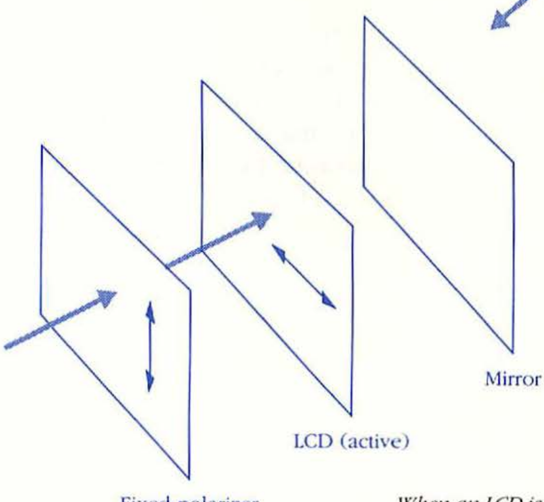

When an LCD is electrically activated, its polarization axis is twisted 90 or 270 degrees with respect to the front fixed polarizer; light is blocked and no longer reflects back to the viewer, who sees a dark spot.

A third type of LCD is an active matrix LCD, which has a transistor at each pixel. When the transistor is on, the pixel stays active continuously, which gives higher contrast than other LCDs.

LCDs consume very little power (as little as 80 milliwatts, compared with about 30 watts for a small CRT), but they are hard to read, especially when the lighting is not optimum. All LCDs—even supertwist and active matrix screens—suffer from both poor contrast and a narrow viewing angle.

Gas plasma displays consist of an array of tiny neon lights that emit a distinctive red glow when electrically excited. Plasma screens give fairly high-contrast images and consume 15 to 20 watts; they are expensive, compared with either CRTs or LCDs.

Electroluminescent (EL) displays are visually the most attractive, featuring high contrast and a wide viewing angle. Their characteristic yellow-green color is created by electrically exciting a mixture of zinc sulfide and manganese. An EL panel consumes as much power as a plasma screen and costs about 20 percent more.

Color flat-panel displays have been built using both LCD and EL technologies. By 1989 such panels may have enough pixels to serve as a Macintosh color screen.

SPEEDING UP DISPLAYS

No matter what the screen, it takes time—usually central processor time—to compute the size, placement, and color of pixels to be displayed. The more pixels and the more colors there are, the longer the computation time and the slower the overall operation. To cut display processing time, you could add an accelerator board to the Mac Plus or SE, switch to a Mac II, or speed up a Mac II.

Another way of dealing with slow display processing is to add a separate video processor to relieve the CPU of computing the display. A video processor can be simply another general-purpose CPU or, better, a processor optimized for video, such as the Texas Instruments 34010. Such chips could be programmed to process the Mac's QuickDraw instructions for creating the screen image or Adobe's PostScript instructions (now used primarily to send printing instructions to a LaserWriter).

Video processors require some changes in the way present software works. Macintosh system software will need modifying so that video commands are directed to the video processor instead of to the CPU. In addition to preprogrammed instructions such as QuickDraw, some applications might send their own commands directly to the video processor. You should be sure that any video processor board you buy works with your software.

LOOKING AT THE SCREEN

Some people find that looking at a computer screen for a long time strains their eyes. Usually, eyestrain can be eased or eliminated by paying careful attention to installation.

If you use the Macintosh while sitting at a table, be sure the table is at typing height—about 29 inches off the floor. At this height, you might find the screen of the small Macs, or a separate Mac II monitor resting directly on the table, a little low. You can raise the Macintosh or monitor, or tilt it back to get a more comfortable viewing angle. If you wear bifocals, tilting will probably give you better results than raising the screen. (Unless you add a fan, be careful not to tilt a Mac Plus or earlier model back so far that you block the ventilation holes on the top; doing so would trap warm air inside the computer.) Several companies make adjustable tilting platforms that fit under the Mac or a display monitor. But since you will probably work at only one angle, a block of wood does just as well and is a lot cheaper.

For a good discussion of computer furniture design, see *Video Display Terminals Handbook* (catalog number 350-049, $8.00 from AT&T Distribution, Indianapolis, IN, 800-432-6600).

Flicker

In average conditions, you should not notice any flicker on the Macintosh screen, although the standard black-and-white display is brighter and therefore more flicker prone than computer displays with a dark background. The problem gets worse the larger the screen, so if you buy a big screen, look for units with higher refresh rates. Because human peripheral vision is especially sensitive to movement, you can see more flicker if you look at the screen out of the corner of your eye. If you notice flicker and find it bothersome, turn down your screen brightness, dim the room lighting slightly, or sit a little farther from the screen.

Glare

The screen surface of most Macintoshes is slightly roughened to reduce reflection, so you shouldn't need to set it up in a darkened room. Simply look for a place out of direct sunlight; the brightness of the screen and the surrounding area should be about the same so that your eyes won't have to keep readjusting.

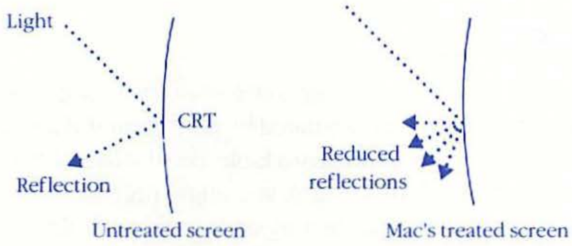

Room lighting often poses problems. Common office lighting with large fluorescent fixtures directly above the work area can make using video displays uncomfortable. Any bright objects opposite the screen—walls or even clothing—can produce a reflection. To suppress such reflections, you might want to put up a dark wall hanging opposite the screen. If you are designing a computer installation, plan the room with indirect lighting or lighting a short distance away from each computer. Desk lamps or lamps on flexible arms near the computer can illuminate work areas without flooding the CRT. Light from large windows can be tamed with screen accessories (see below) or with room dividers, horizontal venetian blinds, or, better yet, vertical louvers that can be adjusted to block the sun and leave you a view.

Still, you might not find a place for your Macintosh that avoids glare completely. If your room has a big picture window, for instance, it would be a shame to give up the view so you could stare at a computer screen. An accessory antiglare screen might help, although some can darken too much.

Mesh screens ($10 to $20) use closely woven, matte-black filaments to scatter room light coming from any direction. Because of the Macintosh's small pixel size, the mesh must be very fine. Sun-flex (Novato, CA) makes the best model of this type. You will probably need to turn up screen brightness, and the image will be a little less sharp. To clean dust from the mesh, blow on it.

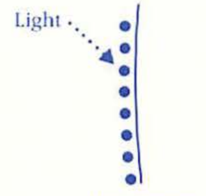

Optically coated screens use antireflection coatings like those on camera lenses and other optical surfaces. The flat, polished glass does not degrade the screen image, but strong light can occasionally create disturbing reflections despite the coating. A small adjustment in screen angle usually fixes this problem. Kensington Microware (New York, NY) sells a hard plastic version.

Interference coatings, which are less expensive ($40), work best in combination with a darkening (neutral-density) filter, which cuts screen brightness.

The best cure for screen glare, however—and the most expensive—is a circular polarizing filter, which dramatically cuts reflection. The best type, made of optical glass, costs more than $100. Suppliers include Vu-Tek (Camarillo, CA) and Polaroid (CP-70). Clean these filters with lens cleaners only.

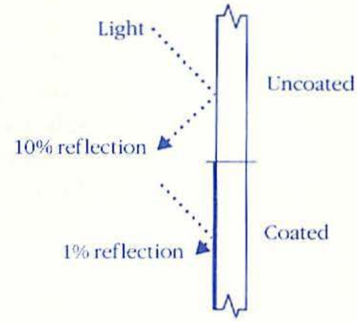

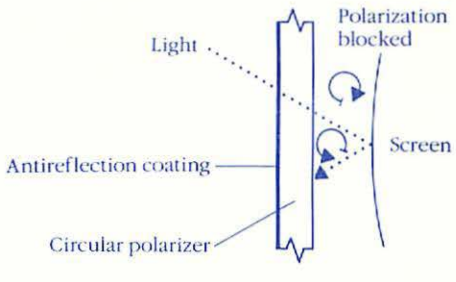

Green and Amber Screens

Some computer companies have touted green or amber screens, supposedly more restful for the eyes, as having special "ergonomic" advantages, but there is actually little evidence one way or the other. If you prefer a green or amber image, green and amber filters designed to fit over the Mac screen are available. If you don't mind the fuzzier image, you could try one of them. For the Mac II, you can select a monitor with green or amber phosphors. All large monochrome screens available so far for the Mac have white phosphors, but colored versions are certainly possible.

TAKING CARE OF THE SCREEN

The Macintosh screen should require little maintenance. If you need to clean your screen, use a household glass cleaner, but try it on a tiny corner first. On some Mac screens, the antireflection roughening is sprayed on, and glass cleaner will simply make a mess; unfortunately, it's hard to tell if a screen's coating has been sprayed on without testing it.

Although you can leave any Macintosh on for extended periods, one problem may develop if you leave the screen illuminated for weeks or months without actively working with it. When an electron beam strikes a screen area continuously, the phosphor there eventually wears. Stationary images left for long periods may thus "burn in," leaving a faint shadow. (You can see this wear on airport television screens that show flight schedules 24 hours a day.) To reduce phosphor wear, you should turn down the brightness control whenever you leave your Mac on and walk away from it for long periods. Many desk accessory programs are available that black out the screen after several minutes' inactivity; a click of the mouse button restores the screen image.

5: The Keyboard and Mouse

With most computers, you enter data and commands from a keyboard. But with the Mac's visual interface, you also need a pointing device—the mouse—to select items, move around the screen, and choose commands.

Apple has now produced four different keyboards for the Macintosh family: the original keyboard, the Mac Plus keyboard, the Apple Keyboard, and the Apple Extended Keyboard. The Apple and Apple Extended keyboards are designed to work with the Mac SE and Mac II. The Apple Extended Keyboard and similar designs from other companies are based on the layout developed by IBM for its personal computers and will accommodate software running under MS-DOS and OS/2, the IBM PC operating systems. Finally, several independent companies are producing keyboards that plug into Macs.

Apple has also produced two different mice: one for Macs before the SE and another for the SE and II.

THE KEYBOARD

Like most English-language typewriters, the Macintosh keyboard has the standard QWERTY layout, named for the first row of letters. A typewriter usually has 84 or 88 characters, including uppercase letters; like most computers, the Mac adds the symbols < > | \ ~ and ^, for a total of 94 characters.

On the Apple and Apple Extended keyboards, the Backspace key has been relabeled Delete because pressing it deletes anything that you have selected on the screen. (This key is not the same as the Del key used by MS-DOS software on the IBM PC or Apple Extended keyboards.)

All the keys on the Macintosh keyboard, except the shift keys, repeat if you hold them down; you can adjust how quickly this happens, as well as the delay before the repeating starts, from the Control Panel desk accessory.

Table 5-1 lists some Macintosh keyboards and their options.

TABLE 5-1. SOME MACINTOSH KEYBOARDS

Keyboard	*Macintosh model*	*Connector*	*Arrow keys*	*Numeric keypad*	*Function keys*	*Escape key*	*Control key*	*Total keys*
Original	128K, 512K	Phone plug	No	No	None	No	No	58
accessory keypad	128K, 512K	Phone plug	L pattern	Yes	None	No	No	18
Mac Plus	Plus	Phone plug	L pattern	Yes	None	No	No	78
DataDesk	All models	Phone plug	Inverted T	Yes	22	Cancel[1]	No	101
Tangent Technologies	512K, Plus	Phone plug	No[2]	Yes	10	Yes	No[3]	84
Apple	SE, II	ADB	Single line	Yes	None	Yes	Yes	81
Apple Extended	SE, II	ADB	Inverted T	Yes	22	Yes	Yes	105

[1]Cancel key functions as an Escape key.
[2]No dedicated arrow keys; numeric keypad keys can work as arrow keys.
[3]Key labeled Control actually works like the Mac Command key.

Special Keys

In addition to the characters and the Shift key (for uppercase letters) and the Caps Lock key (for only uppercase letters), all Macintosh keyboards also have two special shift keys, Option and Command (⌘). You hold down one (or both) of them while pressing another key to get special functions. Keyboards for the Mac SE and the Mac II have an additional special shift key, labeled Control.

Depending on the keyboard, you will also find several other keys, including Enter, Esc (Escape), Clear, cursor (arrow) keys, cursor-control keys, and function keys.

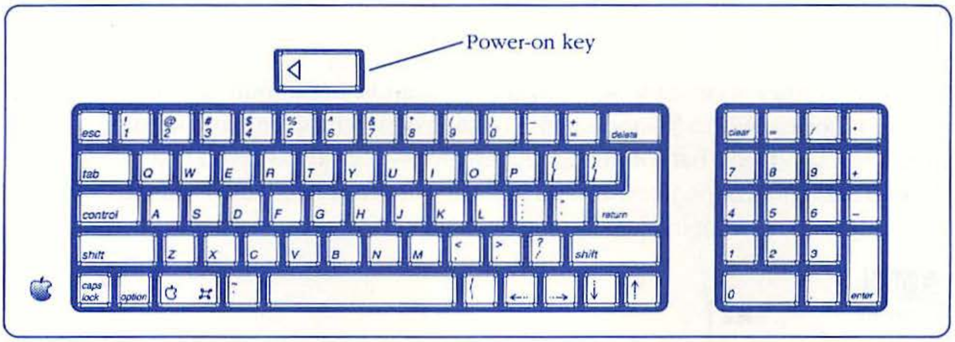

Apple Keyboard.

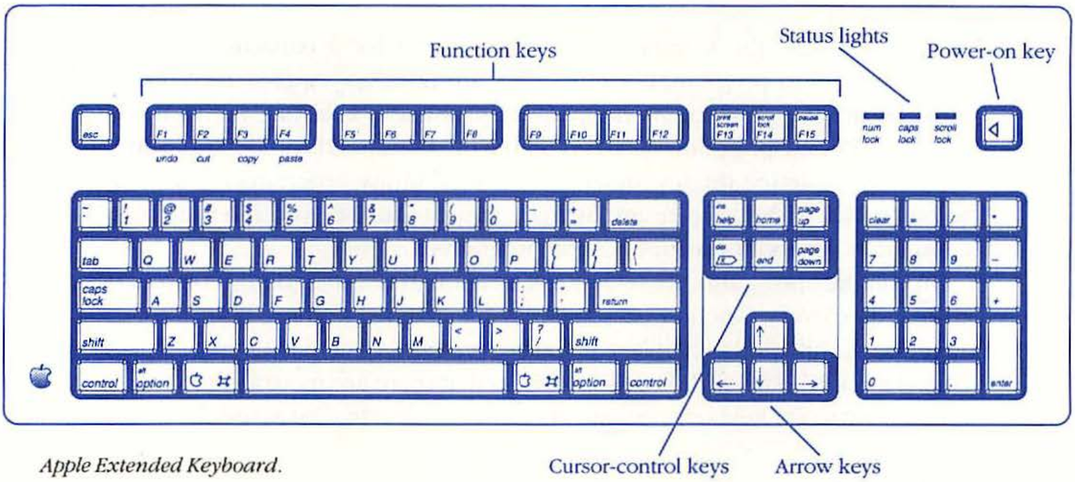

Apple Extended Keyboard.

The Option key

You generally use the Option key to get symbols, special punctuation like dashes, or special characters for a foreign language. The Key Caps desk accessory shows you which characters are available. For MS-DOS programs, the Option key corresponds to the Alt (Alternate) key of IBM PC keyboards. Here are the Option-key characters available in the Chicago font and punctuation marks available in nearly all fonts.

¡ ™ £ ¢ ∞ § ¶ • ª º – ≠
œ ∑ ´ ® † ¥ ¨ ˆ ø π " ' «
å ß ∂ ƒ © □ ∆ □ ¬ … æ
Ω ≈ ç √ ∫ ˜ µ ≤ ≥ ÷

Option-[	“	
Option-Shift-[	”	
Option-]	‘	
Option-Shift-]	’	
Hyphen	-	*(ordinary hyphen)*
Option-hyphen	–	*(en dash)*
Option-Shift-hyphen	—	*(em dash)*

The Command key

The Command key (also called the Apple key on the Apple and Apple Extended keyboards) gives you a shortcut to some menu commands. In the pull-down menus, some commands appear with a ⌘ and a letter; holding down the Command key while pressing the letter produces that command immediately—exactly as if you had pulled down the menu and chosen the item with the mouse. For often-needed commands, you will probably find the Command key faster than using the mouse.

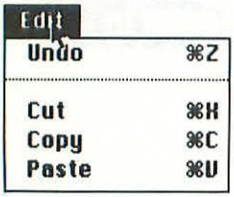

Unfortunately, the Command-key combination for a particular command varies from program to program, but most software does use Apple's recommended combinations for the most common functions (Command-x for Cut, Command-v for Paste, and so on). Some programs let you redefine the Command combinations they use. You can always redefine them yourself with special utility programs or by modifying the resource files within a program. By making such changes, you can standardize certain commands across all your application programs.

In complex programs, there may not be enough Command-key combinations to accommodate every function, so some may be assigned to three keys at a time, such as Command, Shift, and a letter key. If you use programs like these a great deal, you should consider a keyboard with function keys so that there are more keys to work with. You might need to get keyboard enhancer software (see page 76) to use the function keys.

Control characters

When communicating with other computers or running other operating systems, such as MS-DOS or UNIX, you might need to generate control characters. These are characters that perform such functions as ringing the other computer's bell (Control-g) to wake up the operator or telling the other computer when to start a new page (Control-l). The keyboards for the Mac SE and Mac II have a Control key. On other keyboards, the communications program linking your Macintosh with other computers usually makes the Command key the equivalent of a Control key; to produce a Control-c, for example, you type Command-c.

The Esc key

Macintosh software doesn't normally need an Escape key, which is common on other computers and new on Mac SE and Mac II keyboards. If you are using a Mac with an earlier keyboard to communicate with another computer and need to generate an Escape character, your communications program will tell you how to send it. Some programs define the tilde (~) key as Escape (to send a tilde character, you can use Command-~); others simply use Command-[because the code for Escape is equivalent to Control-[.

Arrow keys

Starting with the Mac Plus, Macintosh keyboards incorporate a set of four arrow keys, labeled with arrows pointing in four directions. These keys move the blinking insertion point, or cursor, left or right one character and up or down one line in word processing, between adjacent cells in a spreadsheet, and so on. Like other keys, the arrow keys repeat automatically if you hold them down, and all the shift keys can work with them, as long as your software interprets the combinations.

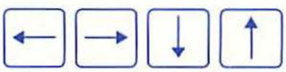

Apple Keyboard

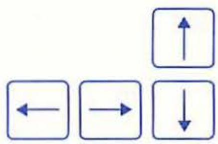

Mac Plus keyboard

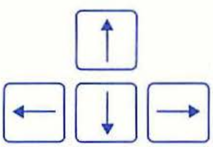

Apple Extended Keyboard

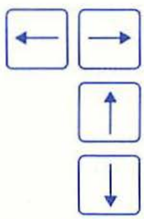

Keypad for 128 KB Mac and Mac 512K

Unfortunately, many Macintosh programs do not make use of these keys. The lack of support for arrow keys stems in part from Apple's decision to leave them out of the original Mac keyboard. When Apple finally did decide to include them, it chose a succession of clumsy layouts—an L shape in the original accessory numeric keypad and on the Mac Plus keyboard and a straight line on the Apple Keyboard for the Mac SE and Mac II. Only on the Extended Keyboard are the arrow keys acceptably arranged in an inverted-T shape.

Arrow keys are excellent for making small, precise movements but are awkward for longer ones; you have to press the keys repeatedly or wait for the automatic repetition to move the insertion point where you want it. For long moves, some programs have set up additional keys or key combinations that move the insertion point a word or a screen at a time. In any case, arrow keys are hopeless for drawing because they give no physical sensation of how far you've moved.

The numeric keypad

Starting with the Mac Plus, all Macintosh keyboards have a built-in numeric keypad. For earlier Macs, there is an accessory keypad, which includes arrow keys. The keys on the numeric pad generate a different key code from the number keys on the main keyboard, enabling software to distinguish the two sets.

The keypad serves two main functions. First, it gives you an adding-machine-like keypad for rapid numerical entries. Most programs that accept numbers will accept input from the numeric keypad. Second, some application programs, such as Microsoft Word, use the keypad to control the insertion point. Pressing the numeric pad 7 key, for example, moves the insertion point to the left edge of whatever line it is on; the numeric pad 1 key moves it to the right edge. Used in combination with Shift, Command, or both, keypad keys select text or move the insertion point by specific increments (by the word, by the line, or to the beginning or end of the document).

The Enter key

The Enter key is located to the right of the space bar on Mac keyboards before the Mac Plus and in the numeric keypad of the Mac Plus and later keyboards. Pressing the Enter key usually marks the end of an entry typed from the keyboard—as when you fill in a blank in a dialog box. In most dialog boxes, pressing the Enter key accepts the default response (generally the one with a double outline).

To give software developers flexibility, Apple has not rigidly defined the Enter key; its function varies from one program to another. In some programs, the Enter key acts like the Return key; in others, it behaves slightly differently. In Microsoft Excel, for example, the selected cell shifts one line down if you press the Return key but remains fixed if you press the Enter key.

If you are running MS-DOS on a Mac, note that the Macintosh Enter key is not the same as the IBM PC Enter key; the PC's Enter key is equivalent to the Mac's Return key.

The Power-on key

The Power-on key on the Apple and Apple Extended keyboards is useful only for the Mac II, where pressing this key turns on the computer. It does not reset or turn off the computer, even though the Mac SE manual calls this the Reset key.

Function keys

In addition to all the keys discussed above, the Apple Extended Keyboard has 21 keys not found on other Macintosh keyboards: 15 of them in a line along the top of the keyboard and 6 in a group above the arrow keys. These are function keys, which can be defined by software to do nearly anything you want them to.

Depending on the particular software, the six keys above the arrow keys allow you to move the insertion point by more than one character, line, or cell at a time. A single keystroke can bring you to the beginning or end of a document, for example, or you can move one screen at a time; some programs give you access to their help utilities. Some Macintosh software makes use of these keys, and they can also be used with MS-DOS software.

No software for the Macintosh is likely to require the use of function keys, because these keys are missing from other Mac keyboards. With some software, using the keys might be optional, and you might be able to define the functions yourself. Keyboard redefinition programs, such as Tempo (Affinity Microsystems, Boulder, CO), also let you define the function keys and other keys as well. Function keys are also helpful for running MS-DOS programs, in which using these keys is customary.

Status Lights

The Apple Extended Keyboard and other keyboards that imitate the IBM PC layout have three status lights at the upper right: Caps Lock, Num Lock, and Scroll Lock. Only the Caps Lock light operates with Macintosh software; it comes on whenever the Caps Lock key is engaged. Currently, the Num (number) Lock and Scroll Lock lights work principally with MS-DOS software.

Which Keyboard to Buy?

Despite its size, the Apple Extended Keyboard is probably the one to get—unless you can find one like it from another manufacturer that feels better. Its width makes getting to the mouse a long reach for right-handers, but this keyboard includes every feature that you're likely to need, even for complex applications. The layout is now standard, used not only with IBM PCs and IBM compatibles but with many other computers as well. Only if you must conserve desk space, and if you are convinced that you will never need function keys or have to work with other computers, should you consider a smaller keyboard.

Connecting the Keyboard

Keyboards connect to the Macintosh in one of two ways: through a telephone-type jack in the Macs before the SE and through the Apple DeskTop Bus (ADB) in the SE and II. ADB connectors are round, with four pins. A variety of pointing devices, as well as conventional keyboards, can be plugged into the computer through these connections. Alternative keyboards are available from DataDesk (Van Nuys, CA) and Tangent (Norcross, GA) for all Macintosh models.

The Macintosh will run whether a keyboard is plugged into it or not. If you leave your machine and don't want anybody using it, you can simply take away the keyboard and mouse. Or on a network, for example, a computer that only stores and distributes files for other computers on the network (a file server) doesn't need a keyboard or mouse. Caution: You might not be able to plug a keyboard or mouse into a Mac with an ADB port when it is already turned on; ADB devices are recognized by the computer only when it starts up.

How the Keyboard Works

The keyboard contains its own processor, ROM, and RAM; the processor scans the keyboard for activity every 3 milliseconds. All but the shift keys are wired in a matrix; the shift keys (Caps Lock, Shift, Option, Command, and Control) are each wired separately. Once a key is pressed, the processor notes the matrix location and whether one or more shift keys have also been pressed. Then a keyboard code is sent on to the CPU, where a program called a keyboard device driver looks up the code and defines the appropriate character. The keyboard has two-key rollover. In other words, if you hold down one key and press a second, the keyboard decodes both keys correctly in sequence.

Keyboard Enhancers

A keyboard enhancer, or redefinition program, lets you assign keystrokes, menu selections, and mouse movements to Command-key combinations and other keys, including function keys. You can assign boilerplate text, such as a return address, to a single key or combination; typing only the key or the combination puts the entire address on the screen. You can also string together a series of keyboard commands, mouse movements, and text, and assign the whole thing to one or two keys. Type those keys, and you can add, for example, "Power to the Revolution!" at the bottom of a letter, save the file on a disk, open a new file, and begin it with "Victory to the Counterrevolution!" Such command sequences are called macros, after computer jargon for short sections of programming code.

The keyboard enhancer Tempo (version 2, pre-release) runs as a desk accessory and enables you to define macros that are program specific. For example, if you define function key 7 to perform one sequence of commands in your spreadsheet, you can also define it to do something entirely different in your word processor; Tempo brings up the correct macro whenever you start one application or the other. To enter the macro, you "record" keystrokes and mouse movements in sequence, which Tempo then stores; once it's recorded, you can then edit it. You can also string together several macros.

QuicKeys, version 1.0 (CE Software, Des Moines, IA) has a variety of common Macintosh procedures built in. You can define a single macro to run in all programs, and you can edit macros, but QuicKeys is not as effective as Tempo for creating long macros. The two programs are compatible and can be run simultaneously if you take care not to overlap keyboard assignments.

AutoMac, version 2.0 (Genesis Micro Software, Bellevue, WA) is a simple keyboard enhancer particularly suited to floppy disk systems because it occupies much less disk space than either Tempo or QuicKeys.

DataDesk (Van Nuys, CA) also has a program that lets you reassign keys, but it works only with the company's 101-key keyboard.

Palm Rests

The early Macintosh keyboards were taller than those made for the SE and II. If you need support for your hands, almost any thin support will do, or you can make one out of wood. Commercial palm rests are also available; choose one with a sloping design.

The Dvorak Keyboard

The standard QWERTY layout used on the Macintosh keyboard was developed by Christopher Scholes in 1873 to slow typists down. The layout, which separated often-used letters from one another, prevented typists from getting ahead of the typebar mechanism in the first Remington typewriter. The QWERTY layout is thus the oldest component of modern computers.

With this century's emphasis on production-line efficiency, many alternative layouts have been proposed to improve typing speed and accuracy. August Dvorak devised the best-known alternative. He analyzed letter frequencies and letter-pair frequencies in English and placed the most-often-used letters where they would be pressed by the most powerful fingers using the least motion. He also analyzed number frequencies, laying down the number row as 7531902468. This rearrangement proved too much even for Dvorak enthusiasts, who have mostly restored the 1234567890 order.

Original Dvorak keyboard arrangement.

Although not optimal perhaps, the Dvorak layout is clearly superior to QWERTY. Speed improves some 5 to 25 percent, and only about half the number of errors occur. The world's typing speed record (170 words per minute) was set on a Dvorak keyboard.

Experiment with Dvorak's design if you're curious. Dvorak layout programs are available for the Mac, but you must change all the caps on the keys, obliterating the standard arrangement. And since even Dvorak advocates admit that only a few thousand Dvorak keyboards exist, choosing Dvorak means choosing isolation.

Dvorak is a lost cause. QWERTY keyboards are like irregular verbs: Everyone hates them but no one can change them.

Keyboard Problems

Apple's keyboards, including those for the Macintosh, have never been outstanding, although many of the problems with layout awkwardness and the lack of important keys, such as arrow keys, have been solved in the Apple Extended Keyboard. Unfortunately, however, keyboard feel still lags behind the precise snap of a genuine IBM PC keyboard—one area where IBM has been and remains convincingly ahead of Apple. So far, independent keyboards for the Mac—like non-IBM-made PC keyboards—are pretty mushy.

The Apple Extended Keyboard and those like it are very wide—so wide that moving your hand from the keys to a mouse on the right side is a minor nuisance. These large keyboards also take up more desk space than other models. (The smallest keyboard that will work on an SE or II is the one sold with the Apple IIGS.)

The ADB connector for the keyboard, as well as the one for the mouse, is located on either side of the new keyboards, instead of on the side facing the computer, as in the early Macs. This means that you need an extra two inches of free space next to your keyboard. Also, if the wire connecting the mouse is stiff, it can be unmanageable. Until right-angle plug adapters are made to fix this problem, you might prefer to connect the mouse directly to the computer rather than to the keyboard.

The proliferation of keyboards for the Macintosh can create problems for fluent typists who regularly switch from one Mac to another—from one at home to another in the office, for example. Here are some keyboard traits to watch out for:

- The main typing area is essentially the same from keyboard to keyboard, although the ' ~ and | \ keys aren't always in the same places.
- The placement of the Caps Lock and Option keys also varies; newer keyboards have added Esc and Control keys, also in varying locations.
- The position and layout of the arrow keys vary greatly.
- Numeric keypads differ in details. (The + and – keys are transposed on the Apple Keyboard with respect to the Mac Plus and Apple Extended keyboards.)

The only way to avoid these incompatibilities is to put identical keyboards on all the machines you use.

Taking Care of the Keyboard

The keyboard will need little, if any, special care. The most common threat is spilled coffee. Many companies have a sensible rule against placing coffee or other drinks on any table occupied by a computer.

You might buy or make a cover to keep dust off. Although many commercial covers are available, you can make a perfectly good one from coated ripstop nylon sold by outdoor and sporting goods stores.

For situations where liquid or debris may fall into the keyboard, such as on a shop floor, flexible transparent keyboard covers are available (SafeSkin from Merritt Computer Products, Dallas, TX). Such covers inhibit touch-typing, but simple keyboard entries are still possible even with the cover in place.

THE MOUSE

You need a keyboard for entering text, but when using the Mac, you are actually pointing at the screen much of the time—choosing a menu item, marking text when word processing, aiming rockets in a game, or drawing.

The mouse remains the most satisfactory general-purpose pointing device. Many studies have shown its superiority, but people are often skeptical until they try one.

Learning to use a mouse takes virtually no time, and once you have worked with it, you will appreciate its rapid, natural operation. You don't have to look at a mouse while operating it, and it stays in place when you let go; your arm generally rests comfortably on the desk. If you run out of desk space in the middle of a movement, you simply pick up the mouse (the pointer stays where you left it on the screen) and go over the same space again to continue your movement.

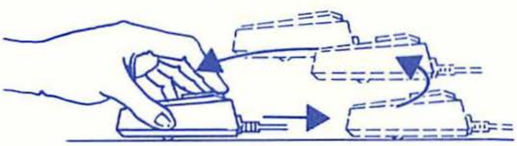

Moving a mouse over the same desk space.

The Apple Mice

Apple has produced two versions of mice, which share the same basic one-button design. The first was produced for the Lisa, early Macintoshes, and the Apple II before the IIGS. It has plastic rollers inside, which pick up dirt and grease quickly, and tiny plastic feet that wear down quickly on a rough surface. The second, which connects via the Apple DeskTop Bus, works with the Mac SE, Mac II, and Apple IIGS. With steel rollers and large Teflon feet, it has a surer feel and a better mechanical design than the original mouse.

The two mice are not compatible or interchangeable.

Disadvantages of a Mouse

A mouse is not ideal for every application. In word processing, for example, using the mouse to move the pointer on screen means taking your hand away from the keyboard, making a movement with the mouse, and moving your hand back to the keyboard again—a particularly time-consuming action if you have a big screen and an Apple Extended Keyboard. The effort is worthwhile for moving a long distance, but small moves are simpler with arrow keys. Unfortunately, most word-processing software for the Mac has not made use of the arrow keys, but this should change now that these keys are standard on Mac keyboards.

The mouse is not ideal for drawing, either, because you can't coordinate your hand and the mouse as well as you can your fingers and a pen.

A mouse requires extra desk space, which may be hard to come by in a crowded office—but then, five years ago you probably didn't have space for a computer. Because it needs room, a mouse is awkward with portable computers as well.

How the Mouse Works

The Macintosh uses a mechanical/optical mouse: The roller on the bottom is mechanically coupled to two rotating vanes with slits. The slits in the vanes interrupt beams from light-emitting diodes (LEDs) that light up phototransistors. As the beams are interrupted, the two vanes track vertical and horizontal motions with respect to the screen.

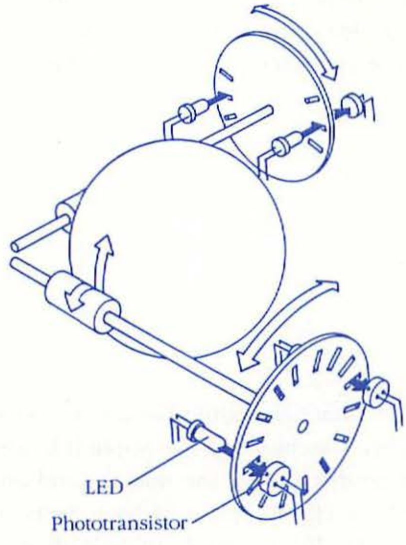

Detecting the direction of motion is a little tricky. The LEDs and phototransistors are in pairs; because of a small offset, the lead or lag of one phototransistor's output with respect to the other gives the direction (in jargon, quadrature modulation).

Vertical and horizontal orientations are defined by the mouse, not your point of view. If you hold the mouse at an angle and move it straight back, the pointer moves vertically on the screen.

Mouse movements depend on the absolute orientation of the mouse.

A mouse registers relative movement; the controller electronics can tell how far the mouse has moved and in which direction, but not the mouse's absolute location. You can adjust tracking speed (how rapidly the pointer moves across the screen when you move the mouse) and double-click timing from the Control Panel desk accessory.

Taking Care of the Mouse

The Macintosh mouse works best on a reasonably smooth, clean surface. A pad on your work table can help keep a rough surface from wearing down the bottom of the mouse. When selecting a mouse pad, be sure its working surface is free of labels or any raised printing that could interfere with mouse tracking. If you don't use a pad, don't bear down hard when you hold the mouse because this can also wear down the bottom.

Keep your mouse area clean because the rollers can get clogged with debris from a dirty tabletop. Even a little dust and oil will make a mouse behave erratically. If yours falls ill, turn it over and undo the ball-retainer ring. The most likely problem is dirty rollers inside the ball chamber. Clean them carefully with a cotton swab or a cloth moistened with alcohol. Then clean the ball, brush and blow out the ball chamber, and reassemble the mouse.

OTHER POINTING DEVICES

Because pointing is so fundamental to computer operation, designers have developed many devices other than the mouse to do it. Most types of pointing devices are now available as Macintosh accessories.

Alternative Mice

Mouse Systems has an optical mouse (the A+ mouse) for both early Macs and the Apple DeskTop Bus. Instead of rollers, infrared LEDs and photosensors track the mouse position over a special pad inscribed with vertical and horizontal grid lines. Optical mice are less likely to become clogged, so they often work better in locations where debris is a problem. The pad can be a nuisance, but many people find they need a pad for a mechanical mouse anyway. Using an optical mouse feels different because it offers much less frictional resistance to motion; it also seems to track more surely. We find the A+ mouse to be superior to any of Apple's mice.

Wireless Mice

The "tail," or wire, connected to a mouse can be bothersome, particularly with Apple's on-the-side plug arrangement on its Apple and Apple Extended keyboards. A wireless mouse operating with radio or infrared signals might be preferable. The mouse could be powered by batteries, which would run down fairly quickly since mice are used continually, or it could be driven by a power coil built into a mouse pad. But a wireless mouse would be more likely to get lost than a conventional one. Although several companies have built prototype wireless mice, none have yet made it into production.

Touch Screens

Touch screens boast the most natural operation of all: You point at something by actually touching it on the screen. Different touch screens use a variety of electronic techniques to sense the location of your finger.

MicroTouch Systems (Woburn, MA) makes touch screens for both the Mac SE and the Mac II. Its glass screen and supporting hardware sense a change in capacitance when you touch the screen.

Touch screens are rare because they don't really work well. They have a place in computer systems that must be simple enough for a passerby to operate, as in a store display, but for most uses they present intractable problems. A finger is too big to select a character on the screen, much less a pixel; few people will want to put their fingers in a pencil sharpener first. And while actually pointing at an object, you can't see it because your finger is in the way. Moreover, fingerprints soil the screen. These problems are bad enough in word processing, but they are untenable for drawing, for which you must have precise pixel-by-pixel control.

Some companies simply offset the cursor so that it will always be half an inch above your finger. You may now be able to see the cursor, but you must still hold your hand up to the screen while jockeying the cursor around or while thinking—a tiring proposition.

Light Pens

A light pen is a pen-shaped device containing a light receptor that is activated when you press the pen against the screen. The receptor detects the electron beam from the CRT as it moves across the screen (see Chapter 4), and a timing circuit uses the beam to locate the pen. Most light pens won't work in the black areas of the screen. The system may have to reverse the video briefly to locate the pen.

Light pens come in several types and work well for some applications, particularly for computer-aided design. They share many of the same problems as touch screens, including arm fatigue: Imagine holding a pen up to a vertical screen while composing your next sentence. Moreover, you must pick up and put down the light pen each time you need it.

A light pen for the small internal Macintosh screens would be difficult to build. You could use one with an external CRT; the pen would require hardware support from the video interface card. As with touch screens, such a pen would work better with larger CRTs. A light-pen-like device might also be used in the future with computers having flat screens that fold down near the desktop.

Joysticks and Trackballs

The joysticks and trackballs so common in arcade games work well for moving laser guns to shoot at attacking monsters, but they lack the rapid, precise control essential for normal computer operation. A few serious computers had trackballs before mice were common, but their use has faded. Unlike a mouse, joysticks and trackballs give you little kinesthetic sense of how far you have moved.

Joysticks have few applications aside from games. Of the three joystick types, only the simplest has appeared so far for the Macintosh; it has four switches (up, down, left, right). The other joystick designs are a continuous-motion stick with two potentiometers—so that the stick stays where you leave it—and the rare strain-gauge variety, which senses how hard you are pushing in any direction.

Internally, a trackball is essentially an upside-down mouse. Trackballs do save table space and aid in rare situations; some people find them easier for drawing smooth, if somewhat uncontrolled, arcs in MacPaint. Several trackballs for the Macintosh have appeared. The key feature to look for are two mouse buttons: One should be momentary contact; the other should be latching. Without a latching switch, you will need two hands to do the equivalent of dragging with a mouse. Trackballs come with either a

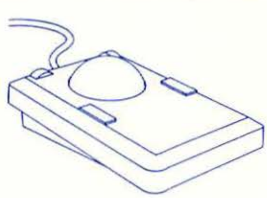

9-pin D-connector (for the early Macs), which plugs in where the mouse does, or with connections for the Apple DeskTop Bus; the two types are not interchangeable.

Joysticks and trackballs can be useful with some physical handicaps; see Chapter 18.

Touch Pads and Digitizer Tablets

Touch pads, operated by a finger, share the resolution problems of touch screens. They also suffer from "skid patches" left by finger oil; movement is often erratic, and fine control is difficult.

Digitizer tablets, which use a stylus, offer excellent resolution and a natural drawing action similar to holding a pen or pencil. Standard tablets work as absolute-position devices—much like a piece of paper and unlike mice. Good digitizer tablets are moderately expensive, and because you must continually pick up and put down the stylus, they do not substitute for a mouse.

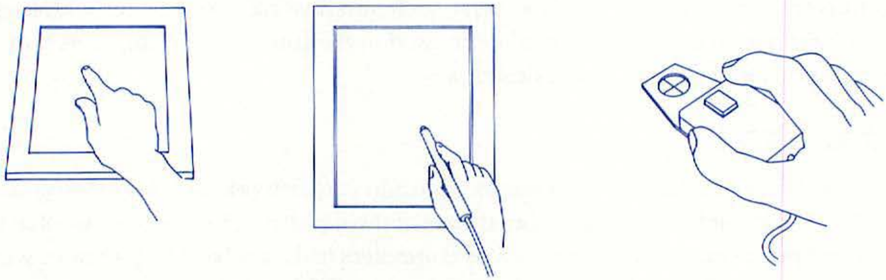

Touch pad; digitizer tablet; and puck, used with a digitizer tablet for tracing.

On the early Macs, the graphics tablets connected through a serial port or through the keyboard port. New graphics tablets for the SE and II connect via the Apple DeskTop Bus. Manufacturers include Summagraphics (MacTablet, Fairfield, CT) and Kurta (Phoenix, AZ). The Kurta pad features an optional wireless stylus that has a tiny radio transmitter. Personal Writer (Los Angeles, CA) has a tablet that can recognize characters printed by hand; it must be trained for an individual user's handwriting.

Additional Pointing Devices

Several other pointing devices can help in special situations: if one hand is occupied with leafing through papers, if you are working on an assembly line, or if you have a physical handicap. A foot-operated, four-way switch can substitute for arrow or other keys. A voice-command system can accept simple spoken instructions.

Lightgate (Emeryville, CA) makes Felix, an unusual pointing device with a stubby handle that you hold like a pencil. The device moves within a one-inch square; small movements are magnified to move the pointer across the screen. Felix has a twitchy feel and is hard to master; it takes tremendous patience to make small, precise movements. Felix does share with trackballs the advantage of needing very little desk space.

Two companies, Stride Computer (Reno, NV) and Position Orientation Systems (Burlington, VT), have developed pointing devices that you direct by moving a small reflector mounted on your head or on your glasses. An infrared light source and a sensor on top of the Macintosh senses the reflector position and angle. Your natural head movements might need some retraining, particularly for side-to-side motions. Although the systems can resolve to the pixel, they are more suitable for text than for drawing.

Stride Computer calls its prototype head-mounted pointing unit the Nod and is seeking general uses for it. Position Orientation Systems is specifically interested in aids for the handicapped; when you hold your head still long enough (the delay is adjustable), this prototype system generates a mouse click. Personics (Concord, MA) is marketing a head-mounted pointing system using ultrasound.

What about the most natural pointing technique of all—tracking where your eye is looking? Eye tracking is extremely difficult. The common systems follow the edge of the iris with photosensors and aren't accurate enough for most computer applications. Precision eye-tracking systems rely on specially designed—and uncomfortable—contact lenses. But even with such equipment, normal eye movements consist of both slow moves and sudden shifts and are too undisciplined for pointing.

6: Disks and Disk Drives

Computers most often use magnetic disk drives and disks for long-term information storage. A disk drive records and plays back digital information much as a cassette recorder records and plays back speech or music, but instead of recording on tape, a disk drive records on a thin plastic or metal disk with a magnetic coating. Disks are reusable; old information can be replaced many times by new information.

Magnetic disks come in two principal types (although intermediates also exist). A floppy disk is a circular sheet of plastic that rotates inside a protective envelope. A hard disk, also called a Winchester disk, is a magnetically coated rigid aluminum or glass platter. Hard disks are much faster than floppy disks and can store much more information; they are indispensable if you use your Macintosh for anything more than simple tasks.

Another medium for mass long-term storage is the optical disc (by convention spelled with a *c*, as distinct from magnetic disks). In optical technology, the equivalent of a disk drive uses a laser beam to read information stored on a rapidly rotating, reflective plastic disc. Optical discs can store far more information per square inch than magnetic media—about five to twenty times as much as a hard disk. Because of their tremendous storage capacity (the entire text of the *Encyclopaedia Britannica* would fit on a single disc less than five inches in diameter), optical storage devices could eventually replace magnetic disks.

MACINTOSH MICROFLOPPY DRIVES AND DISKS

The Macintosh uses $3\frac{1}{2}$-inch microfloppy disks, originally developed by Sony of Japan. Enclosed in a rigid plastic covering, these little disks are far easier to handle than earlier, larger floppy disks.

The standard Macintosh microfloppy disk drive can store nearly 800 KB (800,000 characters) of information on each disk. These so-called double-sided drives have two recording, or read/write, heads; the single-sided drives found on early Macintosh models have only one read/write head and store only 400 KB. Double-sided drives can read and write both double- and single-sided disks, but single-sided drives can read and write only single-sided disks.

New Macintosh high-density disk drives store 1.6 MB on two sides of a high-density disk. For compatibility with early Macs, these drives can still read and write 800 KB and 400 KB microfloppies. In time, the high-density format will replace its predecessors.

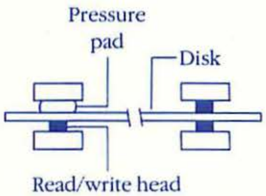

Single-sided disk drives have a pressure pad opposite the read/write head (left); double-sided drives have two heads opposite each other (right).

When you insert a microfloppy disk into a Macintosh disk drive, a lever opens the disk's spring-loaded metal dust cover, exposing the magnetic surface. The two heads of a double-sided drive press on opposite surfaces of the disk. The single head of a single-sided drive reads and writes only on the disk's bottom surface; it presses up while a pressure pad on the other side holds the disk against the head. A motor then moves the head(s) radially across the disk, while another motor turns the disk at 390 to 605 revolutions per minute.

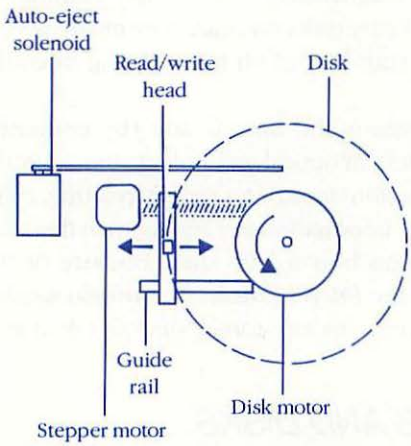

The disk drive mechanism has two motors; one turns the disk, while a stepper motor moves the read/write head radially across the disk.

The read/write head records information in the form of magnetic changes in 80 closed concentric circles called tracks. The outside tracks are longer than the inside ones and can therefore hold more information. To take advantage of this, the Mac's disk-controller circuitry turns the disk slowly when the drive head is at the outer edge and faster when the head is at the inner edge, so the drive head writes more information on the outer tracks than on the inner ones. The changes in motor speed produce the different tones you hear as the drive operates. (Because they operate at a constant rotational speed, most disk drives for computers other than the Mac store the same amount of information in each track, even though the outside tracks have greater capacity; as a result, these drives store less information on each disk.)

Each concentric track on each side of a standard-density Macintosh disk consists of several sectors. The outermost 16 tracks have 12 sectors each; the next 16 tracks contain 11 sectors each; and so on to the innermost 16 tracks, which have 8 sectors each. Each sector contains 512 bytes of information. On a high-density disk, each track contains twice as many sectors as are contained on standard disks, enabling it to store twice as much data. The rate at which data are transferred to and from high-density disks is the same as it is for standard disks, however, because high-density disks rotate in the disk drive at half the speed of standard disks.

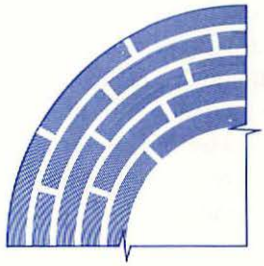

A standard-density Macintosh disk has 80 tracks on a side. The outermost 16 tracks have 12 sectors each; the next 16 tracks have 11 sectors, and so on; the innermost 16 tracks have 8 sectors each.

Buying Disks

Despite their designations, both single- and double-sided disks are magnetically coated on two sides so that the plastic will not warp. During production, the manufacturers test and certify each side. In some cases, double-sided disks that fail on one surface are sold as single-sided disks instead. Because the test and certification process is more stringent than the demands of normal usage, most single-sided disks, which are a little cheaper than double-sided disks, will generally work in a double-sided drive. But rather than risk occasional failure to save a few cents, use certified double-sided disks if you have a double-sided drive. After all, your data are much more valuable than your disks.

Similarly, you should use certified high-density disks in high-density drives. These disks look exactly the same as standard ones, but their magnetic coating is different.

They will be marked "high density" or "2 MB" or both; their raw storage capacity is two megabytes. Even though you can often set up (format, or initialize) a standard disk as a high-density one, performance will be unreliable. In contrast, a high-density disk will work fine in a standard disk drive, although the drive won't be able to write as much information on the disk as the disk can hold. You might want to buy high-density disks even if you don't have a high-density disk drive, because your next computer is likely to use such disks.

Most standard microfloppy disks on the market are specified to record at 135 tracks per inch (tpi), the current standard and the density at which the Macintosh records information. Two-megabyte high-density disks also operate at 135 tpi, although they store twice as much per linear inch of track. In Japan and some other countries, 67.5-tpi disks are also available; avoid them because they tend to fail at Mac's 135 tpi.

All reputable brands of disks work; differences in quality, if any, are hard to discern in normal use. If you are using disks in extreme temperatures or climates, however, do some testing to be sure they are reliable. Always be cautious if you are offered non-brand-name ("generic") disks at an unusually low price, for defective disks do exist.

Write Protection

A common danger when working with any magnetic storage medium is writing over or erasing important information. A small plastic insert called a write-protect tab in one corner of each microfloppy disk lets you guard against this. If you set the insert so that you cannot see it from the disk's top surface, the disk drive, which mechanically senses the write-protect tab, cannot write on the disk. (With some older disks, you have to break off the plastic piece and reinsert it.) On the Macintosh Finder screen, a small icon of a lock appears below the title bar in the window of a write-protected disk.

When it is practical, you should write protect your original program disks and make working copies for daily use. You can't do this with some copy-protected software, however, because you must use the original program disk as a working disk, and the program might require the Mac to write on the program disk as you use it.

If the small square is blocked, you can write on the disk. If the square is open, the disk is write protected.

Initializing Disks

When you buy disks to store your data, they are blank. Before you can use them, the disk drive must initialize, or format, each disk by laying down magnetic markers for each track and sector. If you are reusing a disk, initializing erases any information previously stored on that disk. Some disk suppliers sell preinitialized disks at extra cost; if you buy these, be sure they are initialized for a Mac, not for MS-DOS, OS/2, or another operating system.

If you put in either a brand-new disk or a disk that was used on a computer other than a Macintosh, a utility program in the System file puts a dialog box on the screen asking if you want to initialize the disk and, if so, whether it is to be single- or double-sided (or high-density, if you have such a drive). If you click the Initialize button, the drive initializes the disk; if you don't wish to initialize the disk—for example, if you realize you have put in the wrong disk—you must click the Eject button.

The Disk Directory

How does the Macintosh operating system know where on a disk to find a specific file? Several tracks store the disk's directory, which is an index to its files. The directory, together with a hidden DeskTop file that occupies space on the disk but does not normally appear in the disk window, holds the information—the name of each file and its icon—that you see in the Macintosh Finder. The operating system uses this directory to find a file on the disk.

When you select a file to work with, the operating system first goes to the disk directory, locates the file's name, and notes which tracks the file is stored on. It then moves the disk drive head to those tracks and begins reading the file into RAM.

When you finish your work and want to save it on a disk, the Macintosh first adds a directory entry, finds tracks available for storage (the tracks won't necessarily be in sequence), and then moves to each track in turn to record the information.

If you have run out of space on a disk, an alert box appears on the screen. In most cases (depending on the software you're using), you'll have the chance to switch the disk for another one with more storage space.

Speeding Disk Operation

Simply because it's mechanical, a disk drive operates more slowly than electronic RAM. In addition, a disk drive may perform many steps to read or write a file, moving back and forth between the directory tracks and the data tracks, which may be scattered in several places on the disk.

To speed operation, the Macintosh keeps some disk-directory information in RAM, in an area called the disk buffer. Most programs write changes from the disk buffer on the disk, either immediately or at regular intervals (this is called flushing the disk buffer), but some programs may not take the time to do this. That's why you can't simply remove a disk from the Macintosh but must use the Finder to eject it; the latest directory information is then written on the disk before ejection. (On non-Macintosh computers, you can easily remove your disk at any time—including the wrong time.)

The speed improvement of a disk buffer comes with some risk, however: A power failure will erase the directory information in it. Because no computer is immune to power outages, your best defense is to make frequent backup copies of files and disks; or you can use uninterruptible power supplies, which are discussed in Chapter 17.

Erasing Files and Disks

To erase a file, drag its icon over the Trash icon so that the Trash icon reverses color; then release the mouse button. The Trash icon bulges to show it contains something, and the DeskTop file on the disk notes that your file is in the Trash. The operating system erases files when disk space is needed for a new file, when you choose Empty Trash from the Special menu, when you start an application, or when you eject the disk the file is on by dragging it to the Trash.

If you make a mistake in sending a file to the Trash, you can recover it easily, but best do it soon: Open the Trash (double-click the Trash icon), and drag the icon of the file you want to recover back to the disk window. The operating system restores your file to active status by updating the DeskTop file.

Under some circumstances, you can even retrieve a disk file after an Empty Trash command. If no other writing activity has occurred on the disk, you can use a utility program to recover the file. If writing activity has occurred, you will need some understanding of disk file structure and a program called a sector reader. If you have a sector reader capable of searching for specific strings of text or other data, you can find whatever fragments of a file remain after other disk activity. Some programs are designed to make recovery of lost data easier. (See Chapter 17 for more information.) File recovery was straightforward with the now largely obsolete Macintosh File System; under the current Hierarchical File System, it is much more difficult.

If you cannot erase a file, it is probably locked; unlock it by clicking the Locked box in the file's Get Info window.

If you simply want to erase many disks quickly, you can use an audio or videotape demagnetizer. The Mac treats a demagnetized disk as if it were completely blank and asks if you want to initialize it. Caution: Do not demagnetize hard-disk cartridges or future high-capacity (5 MB or more) microfloppies because you will destroy essential formatting information recorded by the manufacturer.

Using the Macintosh with Other Floppy Disk Formats

Although the Macintosh disk-drive mechanism looks like the 3½-inch microfloppy drives of other microcomputers, it differs in both electrical and mechanical design. Moreover, the Mac stores information on disks in its own unique way. For these reasons, only microfloppy disk drives designed for the Mac will work with it.

With suitable hardware and software, however, you can read and write disks from other computers. You can read and write Lisa 2 and Apple II microfloppy formats with appropriate software, and Apple's high-density disk drives can read and write microfloppy disks formatted in MS-DOS and OS/2.

Floppy disks come in several sizes. The original, and increasingly rare, 8-inch floppies are used mainly on old office computers. The 5¼-inch minifloppy is the most common disk for microcomputers. Both the 5¼-inch and 8-inch sizes have flexible plastic jackets with openings that expose the disk's fragile magnetic surface and require careful handling; a single fingerprint on the exposed surface or a warped jacket can destroy the disk.

Eight-inch disk drives could be connected to the Mac either with a special controller board or through the SCSI port, but no one has yet produced such hardware because demand is low.

Both the Mac SE and Mac II can accept an external 5¼-inch disk drive for reading and writing disks formatted under certain non-Macintosh operating systems. The 5¼-inch drive can be connected through a special floppy disk controller circuit (as Apple has done), through a controller built into an accessory MS-DOS board (as AST has done), or through the SCSI port (as Dayna has done). All three of these can read and write standard 360 KB MS-DOS disks. Dayna's (Salt Lake City, UT) disk drive box can also accept 1.2 MB 5¼-inch floppy disk drives used by the IBM PC/AT. Given appropriate software, all of these drives should also be able to read and write a variety of other microcomputer disks (as long as these are soft sectored and have 48 tpi), including those from many micros using the CP/M-80 operating system and some dedicated word processors. Dayna's product can also accept a 3½-inch microfloppy disk drive that reads and writes in MS-DOS and OS/2 formats.

Sony's 3½-inch microfloppy design has now defeated its competitors in the small-disk market: a 3-inch design from Hitachi, a 3¼-inch one from Dysan, and a 3.9-inch one from IBM. Several even smaller floppy disk formats have appeared: 2½-inch and 2-inch disks that store 360 KB, designed for portable computers, and a 1.85-inch disk format for analog recording of still video images. If either the 2½-inch or 2-inch disk proves popular, manufacturers will probably build a drive for this size disk that can plug into the Mac. Although the 1.85-inch disk doesn't store digital information, the pictures could be converted by a video digitizer into a bit map that Macintosh graphics programs could use.

Taking Care of Floppy Disks

Although you should treat all disks with reasonable care, the microfloppy's semirigid plastic case and metal dust cover mean that ordinary handling will cause no damage; you can't put a fingerprint on the magnetic surface unless you go to the trouble of holding the spring-loaded cover open. Unlike with older 5¼-inch and 8-inch floppies, you can write on a microfloppy label with a ballpoint pen without damaging the disk. But since a determined assault can ruin even microfloppies, keep them away from small children and pets. And even with the cover, disks may not survive coffee spills or direct hits with cigarette ashes.

If you put a rubber band around your microfloppies, don't wrap it across the dust covers; the band could get caught and expose the magnetic surface. Also, don't wrap the disks too tightly, or the covers will deform.

Any disk can be erased by a sufficiently strong magnetic field—more than 50 oersteds (Oe). Fortunately such fields are rare. A typical office has only a few items likely to damage data: the small magnets commonly used on memo boards, copy stands, paper clip holders, and the area directly underneath the electromechanical ringer of traditional telephones. The original ImageWriter has a magnet in its top cover that can damage disks left on top of it. Other potential problems, such as fluorescent light transformers, are almost always too far away to cause damage.

A small physical separation—three inches or more—is sufficient to eliminate nearly all magnetic threats, so the best insurance against magnetic damage is tidy housekeeping. Set aside specific areas on your desk for disks; don't throw them all over. Store disks in a covered container for protection from dust. Most commercial disk storage boxes are overpriced; file card holders (4-by-6-inch size) and plastic shoe boxes work as well. If magnetic damage does occur, you can reinitialize and reuse the disks.

Despite popular concern, the magnetic metal detectors at airport security checkpoints pose no problem, generating less than 5 Oe in most countries. Neither will X-ray inspection machines damage any computer component or magnetic or optical medium. So don't worry about running everything through the X-ray equipment.

When shipping a disk, wrap it in a plastic bag; the cover and the edges of the case are not sealed, so dust could sneak through. If you use a sturdy shipping envelope, further reinforcement is usually unnecessary. Always make a backup disk before shipping.

If a disk gets cracked, bent, or folded during shipping, you may still be able to salvage your data. The outer shell of each disk is lined with a soft material that will provide some protection if the shell gets broken. Open the shell carefully by breaking the corners away from the shutter with a thin knife. If the disk itself is creased or scratched, forget it. If not, transfer the disk carefully (use lint-free photographic gloves) to an opened but intact shell with lining, tape the shell closed, and insert the unit into a disk drive. Transfer all the files immediately to a fresh disk, and throw away the damaged one.

Occasionally, a disk's protective metal shutter gets bent or even torn off—sometimes when you remove the disk from a disk drive. If the magnetic surface has not been damaged, you can still insert the disk in a drive and read it; make a copy immediately and discard the broken disk.

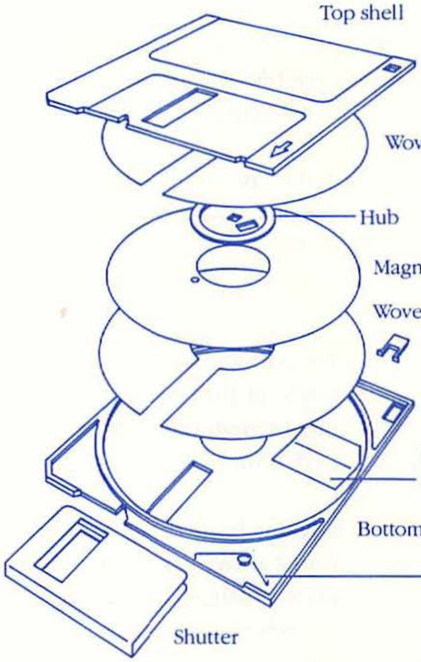

Inside a floppy disk.

Eventually, after a year or more of extended use, the magnetic surface on a disk will wear out. You are only likely to wear out a disk that you use constantly every day. If you have such a disk, make up a fresh working disk after six months and convert the old one to file storage.

Heat can damage floppy disks, so you should not leave disks on heaters or next to the Mac's ventilation slots. Leaving disks in a closed car on a hot day will destroy them.

SHOULD YOU BUY A SECOND DISK DRIVE?

In a word, yes. You can certainly work with only one double-sided disk drive, but a second drive helps tremendously. For maximum performance, the second drive should be a hard disk drive. In some situations—if you take the Mac on the road and knock it around, for example—a second floppy disk drive is better because floppy drives are less fragile than hard disks. Most people with a hard disk drive will have no urgent need for a second floppy disk drive; the exceptions are the few people who do a lot of floppy disk work, such as preparing floppies for distribution within an organization. (If you produce a lot of disks, consider using a disk duplicating service or possibly buying a disk duplicator.)

A third microfloppy disk drive is rarely useful, although an SE can support three—two internal and one external. If you put a third floppy drive on a Mac SE or Mac II, it will probably be a 5¼-inch drive used for transferring files from other computers. Although you could use 5¼-inch disks for storage, a microfloppy stores much more in a smaller package.

If you add an external floppy disk drive to a small Mac, place it to the right of the computer. If you put it on the left, the power supply circuitry inside the Mac will interfere with the operation of the disk drive.

HARD DISK DRIVES

Like floppy disk drives, hard disk drives also record information magnetically and store it in concentric circles on a disk. But the disk is not made of thin, flexible plastic; instead, it is a rigid metal or glass platter coated with magnetic material and polished to a mirror-smooth finish. Hard disk drives are often called Winchester drives, after the name IBM gave its original small-hard-disk technology.

Unlike the heads of a floppy disk drive, which actually touch the disk's magnetic surface during operation, the heads in a hard disk drive fly just above the surface. Even the tiniest speck of dust or cigarette ash will cause the heads to crash, so the platter spins in a chamber of filtered air. Whereas floppy disks spin only when you are reading or writing information on them, hard disks spin continuously. Most hard disk units have two or more platters permanently fixed in a drive.

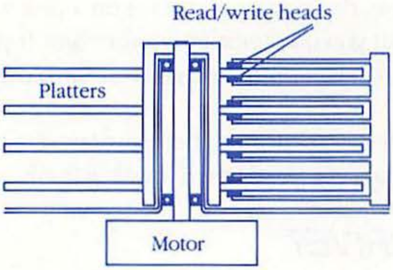

Inside a hard disk drive.

These features mean that hard disks are made much more precisely than floppies and can store much more information in the same amount of space. A microfloppy has 135 tracks per inch; a hard disk has from 400 to more than 900 tpi. The linear recording density of a microfloppy is 5000 bits per inch (bpi); on a hard disk, it is 7000 to more than 40,000 bpi. A micro-hard disk drive, which fits into the same space as the Macintosh microfloppy drive, can store 20 to 120 MB. Other hard disk drives, only a little larger, can store more than 700 MB.

If you regularly use several programs and work with many documents, you will find a hard disk drive nearly essential. Hard drives cost more than floppy drives—about two to ten times as much—but when you consider the cost per character of information stored, they are much cheaper than floppies, and prices have been falling rapidly. Hard drives are more fragile than floppy drives, however, so if your computer gets rough treatment, you should probably stick to floppies. And all hard drives use fans for cooling, which adds noise to a Mac Plus or earlier model.

Connecting a Hard Disk

The standard method for connecting a hard disk to the Mac Plus, SE, or II is via a Small Computer Systems Interface (SCSI) port. (See Chapter 7.) The SCSI port is connected to the Macintosh system bus and can support as many as eight devices at a time. The Mac itself counts as one SCSI device, so you can connect seven others to your system.

SCSI drives can be either internal or external. How and where you connect a hard disk drive depends on which Macintosh model you have.

Internal drives come in different sizes. Internal drives for the Plus and SE must have a 3½-inch format to fit inside the case. Because the Mac Plus case was not originally designed to accept an internal hard disk drive, installing one is a little more complex than it is for the SE or II.

The Mac II can accept a total of three disk drives. Two must be 3½-inch drives; the third can be either a 3½-inch drive or a 5¼-inch half-height drive. (Half-height drives are 42 mm high; full height is 84 mm, the height of the first microcomputer floppy disk drives.) A microfloppy disk normally occupies one 3½-inch position, and a hard disk takes up the 5¼-inch position. The remaining 3½-inch drive position can be either a microfloppy drive or a hard disk drive.

External drives usually come with their own cases, power supplies, and cables; these extras mean that they cost more than internal drives. Most external drives can be attached to the Mac Plus, SE, or II without any physical complications. Jasmine (San Francisco, CA) makes a clever external drive that piggybacks on the rear of the Mac Plus.

Although a hard disk drive from a Mac Plus or SE can be moved into a Mac II, this is worth the effort only if the drive has the performance suitable for a II or is intended only for backup duties, where low performance doesn't matter. In many instances, the disk will have to be reformatted because the II uses a different track layout.

With suitable hardware and software, even a hard disk drive installed in another computer can be used by a Macintosh. Compatible Systems (Boulder, CO), for example, makes a SCSI interface board for IBM PCs and clones that can connect directly to a Mac's SCSI port. Such a connection makes any disk drive—hard or floppy—on an IBM PC available to the Mac; you can even start the Mac from the PC's hard disk. TOPS networking software (Sun Microsystems, Berkeley, CA) also links IBM PCs and Macs, enabling you to use a PC's hard disk to store Mac files. (See Chapter 21.)

What Happens When You Start a Mac from a Hard Disk

- When you turn on the Mac, a start-up program in ROM first looks in the primary floppy disk drive for a system disk. If one is present, it will be used to start the computer. If the disk in the primary drive is not a system disk, the start-up program ejects it. If the primary floppy disk drive is empty, the program looks for a system disk in the second floppy disk drive.
- If no floppy disks are present, the start-up program looks for SCSI hard disk drives, starting at the highest SCSI address. (See Chapter 7.) If one is present, the start-up program reads tracks on the hard disk known as boot tracks for information about the characteristics of the disk drive ("driver information"). If you have specified a start-up device in the Control Panel, the computer will go to that device first and load System information. The address of the start-up device is stored in parameter RAM.
- Information about any other SCSI drives connected to the computer is loaded, in turn, into RAM. All the driver information stays in RAM until the computer is turned off or restarted.

Hard disks for the Mac are normally shipped initialized (formatted), containing driver information. They are accompanied by a floppy disk with a program that can rewrite (reinstall) the driver information and reinitialize the hard disk if need be; these two operations can be done independently. For Apple's hard disks, the program is called HD SC Setup. Other companies supply an equivalent program for their hard disk drives; the various programs are not interchangeable.

Hard Disk Capacity

The capacities of hard disk drives have been growing fast—a good thing, too, since the demand for storage capacity has been growing equally fast. The smallest common hard disk stores 20 MB, sufficient for modest needs. Most Mac II users should not consider drives storing less than 40 MB. (As a practical matter, disk drive manufacturers have rarely produced medium- or high-performance drives that handle less than 40 MB.) Eighty-megabyte hard disks are already common, and 160 MB are not unusual. Still higher-capacity drives are already available, but the cost per megabyte of storage goes up rapidly.

Most high-capacity drives for the Macintosh have their own cases and power supplies and must be hooked up externally because they are full height; the Mac II case is only tall enough for half-height drives internally.

The highest-capacity disk drives are 8-inch or 14-inch designs produced for mini- and mainframe computers. Capacities can exceed 2 gigabytes (1 GB = 1024 MB), with prices to match. Some gigabyte-sized drives are already available for the Mac II.

Hard Disk Drive Performance

Two important performance parameters are transfer speed and access time.

Transfer speed, the rate at which information passes between a disk drive and the central processor, is a function of both hardware and software.

With Apple's hard disks, the SCSI port on a Mac II runs at 1.25 MB per second, the one on the SE at about 600 KB per second, and the one on the Plus at about 300 KB per second. SCSI port components are now available that can operate at much higher speeds than any current Mac—up to 5 MB per second. (By contrast, a floppy disk drive transfers data at about 56 KB per second.) As a practical matter, however, data are transferred at these rates for only brief periods. Even a Mac II with several megabytes of RAM will choke if fed information at 1.25 MB per second for long. But in short bursts, high transfer rates help a computer respond quickly to your commands.

In principle, any hard disk drive can be used with any computer, given the proper controller. But how can the same disk drive adjust to different computers' transfer speeds? Disk-controller circuitry and software can slow the data rate: Instead of reading every sector on a track in one revolution of the disk, the drive can skip sectors, so it takes more than one revolution to read them all. This number is the interleave factor, usually expressed as a ratio of the number of revolutions to 1. Fast drives, such as standard Mac II drives, have an interleave factor of 1:1 (the drive reads an entire track in one revolution). The interleave factor of Apple's 20SC hard disk drive for the SE is 2:1 (the drive reads every other sector during one revolution, so it takes two revolutions to read the entire track); that of Apple's hard disk drive for the Mac Plus is 3:1 (the drive reads every third sector during a revolution). The interleave factors of other hard disk drives may be as high as 8:1.

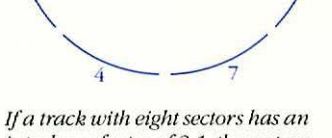

If a track with eight sectors has an interleave factor of 2:1, the sectors are written and read in the order shown. It takes two revolutions to read the entire track.

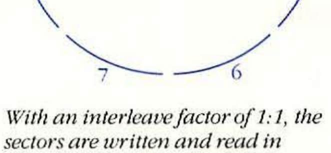

With an interleave factor of 1:1, the sectors are written and read in sequential order. It takes only one revolution to read the entire track.

Average access time is the time it takes a disk drive to locate a particular track on a disk; it takes into account the time to move between adjacent tracks (typically 3 to 5 milliseconds) as well as the time to move to distant tracks. Older hard disk designs typically have average access times of 65 milliseconds or longer. Medium-performance drives take 30–50 milliseconds, fast drives take less than 25 milliseconds, and the fastest take less than 20 milliseconds. In general, the shorter the access time, the more expensive the disk drive.

As with transfer speed, access time depends on the interaction of hardware and software. On the Mac II, a fast drive shines because the Mac II's CPU and bus can handle speedy information transfer. On a Mac Plus or SE, however, speedy disk access can be partially masked by a slower CPU and bus.

Many advertisements for hard disk drives tout misleading speeds, usually by focusing on only one of the many factors affecting speed; fair comparisons are surprisingly hard to do. Even a single drive can vary in how long it takes to read what appear to be identical files. After many files have been created and erased on a disk, the files tend to fragment; the data they contain wind up scattered across the disk and require more drive head movements to read. Many published speed comparisons do not take this fragmentation into account. A utility program, such as PowerUp (Software Power Company, Fremont, CA), can improve a hard disk drive's performance by turning fragmented files into contiguous files on adjacent tracks.

Parking Disk Drive Heads

When a hard disk drive is running, the read/write heads are flying just above the disk surface. What happens when the power goes off? The heads land on the magnetic surface, causing a little wear. A low-capacity hard disk drive can usually withstand such wear for some time, but a high-capacity drive stores so much information on so little surface area that wear can damage data sooner.

Some drives set up one edge of the disk as a "parking" zone where no information is stored. With some drives, you must issue an explicit command to move the heads to that zone. Far better are those with automatic head parking; whenever the power goes off, the heads retract and park. You can often hear the heads retract softly when you cut off the power to these drives. Automatic head parking is an essential feature; do not buy a hard disk drive without it.

Removable Hard Disks

Most hard disks are permanently fixed inside their disk drive, but several designs put the disk in a removable cartridge housing. To get more storage you simply buy another \$50–\$75 cartridge. If you work with unfinished software, a cartridge lets you separate potentially troublesome programs from your important data. The main disadvantages of cartridges include lower storage density, higher price, and, in some cases, unreliability. Cartridge drives have not sold well, in part because prices for conventional hard disks have dropped faster than cartridge drive prices.

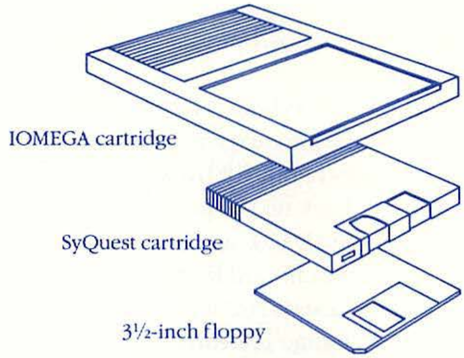

SyQuest produces cartridges and drives packaged and sold by several companies. The Bernoulli Box from Iomega (Roy, UT) and Verbatim's MegaFloppy use a 10 MB cartridge disk that mechanically behaves somewhat like a floppy but performs like a hard disk. Regardless of design, all these cartridges have critical formatting information recorded on them by the factory. Do not use bulk erasers on these cartridges.

Networked Hard Disk Drives

A hard disk drive can be connected to a Macintosh through the AppleTalk network, which connects via the printer serial port instead of through the SCSI port. This allows several people to share files on a network. But a hard disk drive connected through AppleTalk performs slowly, roughly the same as a floppy disk; the peak transfer rate is about 24 KB per second with LocalTalk cabling and can be much slower with multiple users. Faster network performance is possible with other networks, such as EtherNet or enhanced versions of AppleTalk.

A laser printer or typesetting machine can have its own hard disk drive for storing font data or acting as a print spooler, which processes documents for printing while leaving the computer free for other work. (See Chapter 8.) Linotronic typesetters, for example, have a built-in hard disk drive that stores fonts. If installed as a network device, a printer's disk drive could also be available to users for general storage.

Backing Up Hard Disks

With so much information stored on a hard disk, its failure can be devastating. Failures are rare, but for safety you should regularly copy, or back up, your hard disk files. You generally need to copy only those files that have changed since the last backup, so a floppy disk is fine if you create less than 800 KB of new information a day (or 1.6 MB if you have high-density floppies). Floppy disks are inefficient for backing up accounting databases and other large files, but hard disk cartridges are excellent; you will need a combination system—either two cartridge drives or one fixed disk drive and one cartridge. You can also use a network disk server for backups, but someone should back up the server.

Tape cartridges have been the most common backup medium. These cartridges contain long lengths of magnetic tape in cassette shells; they are expensive, slow, and useless for anything except backup, although they are very handy for backing up large disk drives. Some people install a second hard disk just to back up the primary hard disk.

Eventually, a new tape format, Digital Audio Tape (DAT), will replace traditional tape cartridges. The DAT format is a Japanese industry standard for two hours of high-quality digital sound in a cassette a little smaller than a standard audio cassette. A DAT cassette can also store data—about 700 MB. (Raw storage is nearly 2000 MB; redundancy to minimize errors takes up much of the difference.) Any DAT recorder that you use with computers should be designed specifically to work with them; units sold for audio recording will lack necessary controls and a SCSI interface.

Hard Disk Miscellany

Common points about hard disks include the following:

- Space and noise: All external hard disk drives take up space, but they need not sit on your desk. Several designs fit underneath the small Macs; other drives come with cables long enough that they can be tucked away on a shelf (cartridge hard disks are best kept within arm's reach) or even in another room if you have a remote power switch. The advantage of this is that you can get rid of the fan noise.
- Leaving hard disk drives on: Starting up a hard disk drive stresses its motor bearings very slightly, so you should leave it on if you plan to use it again the same day. Although many people leave their drives on for days or weeks at a stretch, Apple recommends turning off a hard disk drive if you won't be using it for eight hours or more. If you leave a hard disk running constantly, the bearings will eventually wear down.
- Mixing hard disks from different vendors: Usually you can use two different brands of SCSI hard disks together with the same Macintosh. Check with the vendors if there is any question. Networks such as AppleTalk support multiple hard disks.
- File recovery: If a hard disk drive fails ("crashes"), the manufacturer can often repair it and recover your data. If you have stored sensitive information on your disk, however, think twice before sending it out for repair; your information goes out with the disk. With the proper software tools, you can do some types of recovery yourself. (See Chapter 17.)
- Plated media: The traditional magnetic coating on hard disks is iron oxide in a binder; a plated media surface consists instead of a thin coating of a metal alloy. Such surfaces are a little harder and tougher, and because they have smaller magnetic particles, they can support higher-density storage.

- Run-length-limited (RLL) coding: RLL is a particular coding technique a hard disk controller uses to store information on a hard disk. RLL operates a little like shorthand, coding for multiple bits when possible. It can typically store 50 percent more information than conventional bit-by-bit coding. Although RLL has been used for many years, particularly in high-capacity hard disks, marketing departments only "discovered" RLL in 1987 and began including the term in advertisements. From the user's standpoint, whether a hard disk drive uses RLL coding is less important than the disk's overall storage capacity. Although some conventionally coded hard disks can be adapted to RLL to increase their capacity, the change requires a new disk controller and careful testing. Most people should simply use their hard disk drive the way it comes out of the box.
- Whitney disk drives: A Whitney drive is an improvement on the Winchester drive because it has a much smaller read/write head, and a smaller head can write on more tracks by going closer to the edges of the disk. A Whitney head is also thinner, so more disk platters can be stacked together in a drive of a given size. But again, whether a particular drive uses Winchester or Whitney technology is less important than overall performance.

Hard Disks for Older Macs

If you need to put a hard disk drive on a Mac that lacks a SCSI port (a 128 KB Mac, a Mac 512K, or a Mac 512KE), you have several alternatives. By and large, you will be better off selling the old Mac and replacing it with a newer model than adding a hard drive to it. In decreasing order of desirability, you could:

- Replace your Mac with a Mac SE or II.
- Replace your Mac with a Mac Plus.
- Add a SCSI port: The port should come from the same vendor as the hard disk drive, especially since not all add-on SCSI ports are the same; this way, if you have any problems, you can go to a single company. A SCSI port can be most easily added to the Mac 512KE, which has routines for supporting such a port in its ROM. Performance is acceptable.
- Connect a hard disk drive through the floppy disk drive port: The principal hard disk drive that works in this fashion is the HD20 from Apple. Performance is barely acceptable.
- Connect a hard disk drive through a serial port: Disk drives you can connect this way are now mostly out of production, so you will have to find a used one. Examine it very carefully. Check with a user group to see if the particular model you are considering has performed well. Do not buy brands from companies that are no longer in business or models that are not widely known to be reliable. Serial-port hard disk drives, whether new or used, should be cheap because they are obsolete; certainly they should not cost more than $200. Performance is barely acceptable.

- Try to connect a hard disk drive internally: Drives that could be hooked up internally to early Macs were first popularized by General Computer with its Hyperdrive. These drives require modifying the Mac, and reliability has been a major problem. Even though some people have used these hard disk drives successfully, you're better off staying away from them. When they work, however, performance is acceptable.

OPTICAL DISCS

Although all optical storage systems use a laser to retrieve information stored on a reflective disc, each system at present requires its own drives and discs; drives of one type cannot read or write discs of any other type. Some interchangeability should be available in the future.

Optical storage systems fall into three main classes:

- Read-only systems: The discs can be read but not written on. CD-ROM (compact disc read-only memory), derived from audio compact disc technology and available now, is the main format.
- WORM systems (write once, read many): The discs can be written on, but only once in a given area; the information can be read as often as needed.
- Erasable systems: You can write and read anywhere on the disc as often as you like, as with magnetic systems. In 1988, erasable discs were largely confined to laboratories, although some experimental systems were available for sale.

Most optical disc drives connect to a Macintosh via the SCSI port, although some may come with their own controller card to plug into the Mac II. Optical drives will not fit inside any current Mac; the Mac II has the physical space—if you're willing to give up the main hard disk drive position—but there is no external opening for inserting an optical disc. A replacement cover for the Mac II could make installation of an optical drive practical.

CD-ROM

Because it takes advantage of the mass production of audio CD players and disc-manufacturing plants, CD-ROM is the most highly developed optical data-storage system. The drive mechanism is the same as that of an audio CD player; many CD-ROM drives will, in fact, also play audio CDs. Audio CD players will not usually work as CD-ROM drives, however, because they are designed solely to convert digital information into analog audio form for a hi-fi system.

CD-ROM discs cannot be written on, so the technology will mainly serve as a publication medium. The discs are mass produced by a stamping process that duplicates an entire disc in a few seconds. Once a master disc is made, each reproduction costs only a dollar or two to press; CD-ROM is the cheapest publishing medium by far, much cheaper than printing books. Nevertheless, most early CD-ROM products were highly specialized databases that cost thousands of dollars—a complete medical index of poisons or indexes to scientific publications, for example. These databases were targeted at relatively small markets willing to pay the high cost. Products that appeal to a wide market, such as collections of general reference works or telephone directories, are beginning to appear.

Compared with magnetic disk drives, CD-ROM players are slow. Access times range from several hundred milliseconds to more than a second, versus 20–65 milliseconds for a hard disk. Faster CD-ROM players are possible, but costs will be much higher because they must use components different from mass-produced audio CD players.

A 4.7-inch-diameter CD-ROM disc contains a single spiral track 5 kilometers long. Information is stored as tiny pits in this track. As the CD-ROM drive shines its laser beam along the track, a photosensor monitors the beam's reflection off the disc. Whenever the beam strikes a pit, the reflection is disrupted, enabling the drive to reconstruct the information.

CD-ROMs, like other computer storage media (magnetic or optical), require formatting, or initializing, to enable drives to find the stored information. There are two levels of formatting: low and high. A low-level format defines fundamental parameters, such as track density or the order of bits on a disc. For CD-ROM the low-level format is called High Sierra, after the location where a meeting to set the standard for the format took place. (A slightly revised version of the High Sierra format is called ISO 9660.) This format ensures that a CD-ROM disc can be physically read by different brands of CD-ROM players, in somewhat the same manner as a standard groove size lets you play a record on any brand of record player. For CD-ROM, as for magnetic disks, high-level formatting is handled by the computer's operating system; high-level formatting sets up disc directories and defines where a particular file begins and ends.

It takes more than correct formatting to get useful information out of a particular CD-ROM. You also need a matched retrieval program in the computer. So far, this software has been distributed on magnetic disk along with each CD-ROM. Eventually, the industry may agree on standard retrieval software for at least some classes of CD-ROMs—books on disc, for example.

In many cases, a file on a CD-ROM disc can be read by both IBM PCs and Macintoshes, provided that suitable retrieval software is available for the respective machines. This software will be able to read a disc format across operating systems. Some applications, however, such as those involving fast-moving animation and digitized sound, might require separate discs for IBM PCs and Macs because of differences in the order in which Intel (IBM PC) and Motorola (Macintosh) central processor chips accept bytes for processing. Only when data must be processed very rapidly do these differences matter; in most situations, the CPU can simply transpose the bytes before processing. Or if the data stored on a CD take up less than half the disc capacity—which will be common—a single disc could contain both IBM and Mac versions.

The raw storage capacity of a CD-ROM disc is nearly a gigabyte. This entire capacity is not available for storing data, however, because some space is taken up by formatting and a built-in error-correction system. Even a tiny flaw in any pit on an optical disc produces an error during reading. Despite careful processing, CD-ROMs have a higher raw error rate than magnetic disks have. To compensate, the error-correction system stores data in partially redundant form. The larger the error (the greater the number of flawed pits), the more redundancy is required. Robust error correction thus both eats up disc space and slows playback of usable information.

CD-ROMs come in two modes with different levels of error correction. Mode 1, which will probably become the most common, holds 556 MB with robust error correction and plays back at 150 KB per second. Mode 2 has simpler error correction, storing 630 MB on a disc and playing back at 171 KB per second. Its greater capacity and faster playback make mode 2 preferable for audio and fast-moving animation, in which an occasional error isn't a problem. For financial data or other precise information, mode 1 is essential.

Because CD-ROMs store so much information, most discs require an index to their data, just as large books need indexes. A comprehensive index may occupy as much storage space as the material it indexes.

Virtually all CDs made so far, whether CD-ROMs or audio CDs, have been single sided. A double-sided CD could be made by pressing both sides or by gluing two discs back to back, but unless double-sided CD-ROM players with two reading heads become available, you would have to turn the disc over to read both sides. Producing double-sided discs poses a practical problem as well. Some fraction of any disc production run is faulty; making two-sided discs would increase the proportion of defective discs. A less risky solution is to build a CD changer. Jukebox systems now being developed can handle from six to hundreds of discs.

WORM Drives

Despite its attractions, CD-ROM technology suffers from a major drawback: You can't write on the discs. You can write on a WORM disc, albeit only once. Yet for many jobs, once is better than not at all. In other situations, the inability to erase can even be an advantage; WORM recording of accounting information, for example, would ensure an audit trail, since the disc could not be erased by accident. Another advantage of WORM discs is their relatively high capacity, which permits you to use them as substitutes for magnetic disks if you create files sparingly. When a disk is full, you simply put it in archival storage.

WORM discs follow the same general principles as CD-ROM but do not require expensive specialized equipment to make and replicate a master disc. WORM drives contain both recording and playback hardware. A typical WORM disc is coated with a thermoplastic, which deforms when heated. During writing, the laser beam operates at high power, heating up and deforming a small bubble on the coating. During playback, the laser operates at lower power, scanning for the deformations and recovering them as data; at this lower power, the beam does not deform the coating.

Half a dozen companies make WORM systems that are available now. Small drives sold mostly for microcomputers store 200 to 400 MB on a 5-inch disc; larger versions can store as much as 3 GB on a 14-inch disc. WORMs can perform much better than CD-ROM; some manufacturers claim access times of 100 milliseconds.

Like CD-ROM, a WORM disc can be double sided, but unless the drive's read/write mechanism is duplicated, the disc must be turned over to use both sides. Some WORM makers also produce jukebox versions.

WORM systems are relatively expensive: A WORM drive costs more than a CD-ROM drive, and each blank disc costs $50–$200. Although databases could be published on WORM discs, the discs must be recorded track by track, a process that takes anywhere from many minutes to more than an hour, compared with the few seconds it takes to press a CD-ROM disc. Several companies are investigating WORM drives that can also read CD-ROM discs or even write a CD (once); the prices for these have not yet been established.

Erasable Discs

Erasable discs involve more complex technology than either CD-ROM or WORM. Two major types of erasable discs are under development: thermo-magneto-optic (TMO) and phase change.

TMO discs are coated with an alloy that holds a magnetic field and whose coercivity, or resistance to a change in the orientation of the magnetic field, is reduced when the alloy is heated. To record information, a laser beam heats the alloy, and a small external magnetic field orients the heated spot. The rest of the disc is cooler and unaffected by the magnetic field. For reading, a low-intensity polarized laser beam is focused on the spot. The orientation of the magnetic field—up or down—slightly rotates the beam's polarization; a sensor analyzes the change to recover the data.

The operation of a phase-change disc depends on switching the state of a coating between a highly reflective crystalline (organized) state and a low-reflectivity amorphous (disorganized) state. To record information, a short, high-powered laser pulse with rapid cooling converts the coating from its crystalline to its amorphous state; to erase the information, a longer, low-powered laser beam with slow cooling reverses the process. Another form of phase-change disc uses a crystal alloy that shifts color when heated or cooled.

In 1988, erasable discs began to come out of the laboratory but were not yet ready for large-scale production. The prices for erasable drives and media will certainly be high initially; their ultimate cost is not yet known.

7: Data Pathways and Memory

The heart of any computer's ability to work with information is the hardware that transfers the information in and out and moves it within the computer. Ports are the electrical pathways that carry information in and out; they are often called I/O ports (for input/output). The bus shuttles information among memory (the computer's workspace); the central microprocessor; video display boards; disk drives; and any accessories, or peripherals, plugged into the ports. Unlike in early Macintosh models, the buses in the Mac SE and Mac II can be expanded to accommodate a variety of accessories.

MACINTOSH PORTS

Ports serve many purposes, handling everything from low-speed communication with a teletype machine to high-speed communication with a video screen. Today's Macintoshes have a variety of ports, both low and high speed.

Based on how they work, ports can be classified in several ways. Each port conforms to a particular protocol, a set of rules specifying the timing, wiring, and voltage that a computer and its accessories must follow for communication. One fundamental distinction, for example, separates parallel from serial ports.

Parallel ports transmit or receive all eight bits in a byte simultaneously on eight separate wires; they cannot transmit and receive at the same time. Serial ports, in contrast, send the eight bits one after another on one wire and receive on a second wire; they can send and receive simultaneously. In general, parallel ports operate faster than serial ports, but the speed difference has little significance for many uses. Ordinary printers, for instance, work so much more slowly than computers that a serial connection is perfectly adequate.

By convention, in common computer usage the term *serial port* refers only to ports like the Macintosh printer and modem ports, not to other kinds of ports that also transmit data in series, such as keyboard ports. Similarly, *parallel port* refers to a specific type of printer port (often called a Centronics parallel port, after the first company to make it). Although the Macintoshes do not have such a parallel port, several companies make serial-to-parallel converters for connecting parallel printers.

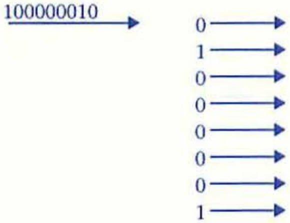

A serial link (left) has a single wire; a parallel link (right) has eight wires. Both links have an additional ground wire; for immunity to electrical noise, the Macintosh serial ports can use two signal wires plus a ground.

The original 128 KB Macintosh, the Mac 512K, and the Mac 512KE come equipped with the following:

- Two serial ports, one for a printer and one for a modem (a device that connects computers to the telephone line).
- A floppy disk drive port.
- Separate keyboard and mouse ports.
- An audio output port.

In addition to these, the Mac Plus and SE also have:

- A Small Computer System Interface (SCSI) port.

The Mac II also has a SCSI port, but it does not have a floppy disk drive port. In place of the separate keyboard and mouse ports, the Macintosh SE and Mac II feature:

- Two Apple DeskTop Bus (ADB) ports.

Because they accept expansion cards, the Mac SE and Mac II can be fitted with many other types of ports for specific applications. (See Chapter 18; for technical information on ports generally, including wiring, see the Appendix.)

The Serial Ports

All Macintoshes have two serial ports that take information both to and from the computer. The three most common uses for these serial ports are to connect a modem, to connect a printer, and to hook up to the AppleTalk network; you can, however, connect many other devices. You can also link two Macintoshes together directly through their serial ports (without using AppleTalk), or you can hook up with an IBM PC or other computer. And you can connect a MIDI interface for electronic musical instruments. (See Chapter 18.)

On the back of the Macintosh, icons identify one serial port as a printer port and the other as a modem port, but the two are functionally the same for many purposes. On the small Macs, the printer port should be used for AppleTalk; on the Mac II, either port can connect to AppleTalk. On the Mac II, the modem port—but not the printer port—has a signal connector for synchronous modems, which are mostly used for communicating with a mainframe computer; the far more common asynchronous modems can be connected to either port.

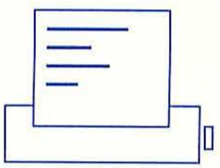

Printer port icon.

Modem port icon.

Macintosh ports follow the RS-422 protocol, an enhanced form of an older, more common protocol called RS-232C. RS-422 improvements over RS-232C allow higher-speed communication and longer connecting wires between devices. RS-232C links are generally limited to 19,200 bits per second (2.4 KB per second), whereas the Mac's RS-422 can run routinely at 230,400 bits per second (29 KB per second) and as fast as 920,000 bits per second (115 KB per second) with special techniques. RS-422 is also much less affected by electrical noise than RS-232C.

On the 128 KB Mac and on the Mac 512K and 512KE, the serial ports use 9-pin D-connectors and supply 5- and 12-volt power to equipment attached to them. For the Mac Plus, SE, and II, Apple changed to round 8-pin plugs (Mini-8 or Mini-DIN connectors) and omitted the power. Some equipment, such as the ThunderScan scanner, which relies on the serial port power supply, therefore needs a power supply adapter to work with later Macs.

The settings for the serial port operating parameters are stored in parameter RAM—which is backed up by the battery—on the clock/calendar chip. (See Chapter 24.)

Connecting accessories

The Macintosh serial ports will handle most RS-232C devices as well as RS-422 devices. RS-232C communications operate at higher voltages than RS-422, but these voltages will not damage the Macintosh serial ports. In many cases, an RS-232C device can be connected with nothing more than a suitable cable; the ImageWriter, for example, is actually an RS-232C device. Most ordinary RS-232C devices sold for microcomputers should work with the Mac, as long as supporting software is available and you can figure out how to set the hardware and software. If the device is not sold for Macintoshes, you may find that getting it to work properly can be a major chore; serial port protocols are often difficult to set.

If you use more than two serial port accessories—more than a printer and a modem, say—you'll need more than two serial ports. Many devices, ranging from television cameras to plotters, connect through these ports. On the Mac SE and II, you could add more serial ports by adding an accessory card, provided supporting software can make the additional ports do what you need.

If you can't add more ports, you can install a switch for selecting among two or more accessories connected to a single port. But if you do, be careful about when you use it. Never flip a switch while data are passing through it, and be sure the switch is set to the accessory you want before sending that accessory any information from the computer. Many devices require an initializing code to begin operation; if you switch to a device after the computer has sent the code, the device won't respond.

Many manufacturers warn you against switching between accessories while the power is on to an accessory, the computer, or both—in part because they don't know what you have connected their devices to. Although many computer peripherals can be plugged into serial ports while the power is on without damaging any components, it is always safest to turn off the power before connecting anything. When in doubt, follow the manufacturer's recommendations.

Several companies make passive switching boxes, ranging in cost from $35 to $100. These boxes contain no circuitry, only switches. Connectors and switches are among the least reliable computer components, so construction quality matters here. The boxes vary in detail. The best can switch eight or nine wires (pins) and connect devices at some distance from the computer or one another via long runs of cable. Less expensive boxes switch fewer pins and are adequate for short-distance cabling.

MacEnhancer, sold by Phoenix Technology's peripheral systems division (Honolulu, HI), is a software-selected switching box that you plug into one of the Mac's serial ports; instead of turning a knob, you set the switch from a desk accessory program.

MacEnhancer gives you four ways to connect. The first is an RS-422 port that replaces the port MacEnhancer plugs into. Two are IBM PC-compatible RS-232C ports (wired according to the DTE [data terminal equipment] configuration) suitable for many devices, including most low-speed devices that would otherwise connect directly to the Mac. The fourth connector is an IBM PC-compatible Centronics parallel port; internal circuitry performs the serial-to-parallel conversion.

The Floppy Disk Drive Port

The floppy disk drive port is bidirectional and operates serially at 62.5 KB per second. It is designed to support one external microfloppy disk drive (in addition to the internal drive or drives). On pre-SE Macintoshes, a second floppy disk drive could only be attached externally through this port. In the original Mac and in the Mac 512K and 512KE, which do not have a SCSI port, the floppy disk drive port can also accept a hard disk drive, such as the Apple HD 20.

The SCSI Port

The SCSI port built into the Mac Plus, SE, and II is a high-speed, bidirectional, parallel port most often used for connecting a hard disk drive; it is not compatible with a Centronics parallel port. The SCSI port can support multiple hard disks as well as a variety of other devices, including tape drives, optical disc drives, fast printers, interfaces to large computers, and small local area networks. A maximum of eight devices, including the Macintosh itself, can be connected one after the other through the SCSI port; this practice of wiring devices together in sequence so that they can all be connected at the same time is called daisy-chaining. At a given time, however, only one SCSI device can actually be sending data.

Each SCSI device has a distinct address, from 0 through 7, usually set with a switch on the device itself. The Mac is always at address 7; in the Mac SE and II, the internal hard disk, which is plugged into an internal connector that is daisy-chained to the SCSI port, is always at address 6. The higher the address, the higher the priority in sending data.

If you have multiple SCSI devices connected to your Mac, be sure their addresses do not conflict. You can change the address of most SCSI devices by resetting their switches, although you might need to have the manufacturer do so. The best-designed devices have an external switch that you can reset easily; watch out for push-button switches, which can be reset inadvertently.

SCSI is a standard interface whose characteristics have been defined by a committee of the American National Standards Institute (ANSI X3T9.2). Many computers, including IBM PCs with appropriate hardware, support SCSI devices. To use a SCSI device on any computer, you will need suitable driver software—software in the computer's system file or operating system that manages information between the computer and the device.

Apple has conformed to the ANSI standard definition of SCSI, with two exceptions. Standard SCSI connectors are 50-pin ribbon connectors, but to save space, Apple has used 25-pin connectors (DB-25). Apple has also not provided any termination resistors, which supply the correct electrical features for ending a series of SCSI devices.

To connect a standard SCSI device to a Macintosh SCSI port, you must use a 25-to-50-pin adapter cable. The DB-25 connectors are physically identical to the plugs used by many RS-232C devices. Do not confuse the two; *never* connect an RS-232C device to the SCSI port. An RS-232C device can put out signals as high as ±25 volts; the SCSI interface is designed to handle only ±5 volts and can be damaged by anything higher. The circuits can also be damaged by static discharges.

Most SCSI devices have two standard 50-pin connectors. With a 25-to-50-pin adapter cable, you can plug either into the Mac; the other one can be daisy-chained to another SCSI device. The last SCSI device in the daisy chain might need a termination resistor; see your hardware manuals.

The Apple DeskTop Bus Port

The Apple DeskTop Bus (ADB) is a low-speed port designed mainly for input devices such as a keyboard and a mouse; the maximum data transfer rate over ADB is 4.5 KB a second. Besides pointing devices from trackballs to graphics tablets, ADB can also support bar-code readers, simple optical-character-recognition machines, and more. ADB is too slow to support sophisticated musical keyboards, however, although it will handle a simple keyboard that plays only one note at a time.

ADB functions like a local area network: Each device on the bus has its own (simple) processor, and information travels to and from each device. The Mac needs device driver software for each device running on ADB. Either the driver or the device can initiate data transfer on the bus. Because device drivers are loaded into RAM when the computer starts up, a device connected after the computer is already on might not be recognized, so you should avoid connecting devices to ADB while the Mac is running. The Mac SE and II have two built-in device drivers, one for a mouse and one for the keyboard, which must be installed in your System Folder.

ADB allows for 16 addresses and can thus support a maximum of 16 devices at a time. But many devices can be designated to share the same address; the software device drivers can distinguish among them by comparing the contents of the memory registers built into every ADB device.

The Mac SE and II have two equivalent ADB connectors in the main case; the Apple and Apple Extended keyboards also have two ADB connectors. The mouse can be plugged into either the main case or the keyboard, or it can be daisy-chained to any other device connected to ADB that has a free socket.

One wire in ADB cabling is used only in the Mac II. It carries a power-on signal that enables you to turn on the computer from the keyboard. The power-on signal could also be generated by another device, such as a modem, making it possible to turn on a Mac II from another machine somewhere else.

The Apple IIGS microcomputer also uses ADB for its keyboard and mouse; ADB devices can be interchanged between the IIGS and Macs, provided that suitable device drivers are available.

The Keyboard and Mouse Ports

Instead of the Apple DeskTop Bus, Macintoshes produced before the SE have separate keyboard and mouse ports. The keyboard port is a slow, general-purpose bidirectional serial port, able to support eight daisy-chained devices, including the optional numeric keypad. Olduvai Software (South Miami, FL) has announced a hardware accessory that adapts the keyboard port of pre-SE Macs to ADB.

The mouse port is suitable only for attaching a mouse or other pointing device, such as a trackball or joystick, although these too can be daisy-chained.

The Audio Output Port

The audio output port transmits audio output only; it is normally connected to a built-in speaker whose volume can usually be set from the Control Panel desk accessory. When set at minimum, the speaker should be inaudible. If you want to silence the speaker at all times so that even the power-on tone doesn't sound, insert a dummy jack with no connections (Radio Shack #274-286 for the small Macs and #274-284 for the Mac II) into the audio output port.

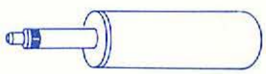

Plug for the audio output port (Radio Shack #274-286).

The audio output port lets you connect Macintosh sound output to a hi-fi system or tape recorder. On the small Macs, the socket is a standard 1/8-inch monophonic minijack. The signal can drive a small speaker and should work satisfactorily through any line-level input on a hi-fi amplifier or receiver (tuner, tape, and auxiliary, but *not* magnetic phono or microphone inputs).

The Mac II audio output socket is a standard 1/8-inch stereo minijack (tip = left channel, ring = right channel, sleeve = ground); it produces a signal at line level (1 volt peak to peak, with an output impedance of 47 ohms). Like the jack in the small Macs, the Mac II jack can be connected to any line-level input of a stereo amplifier. If the audio output port is not connected to anything, monophonic sound combining both channels comes through the built-in speaker.

If you try to connect any Mac's audio output port to a magnetic phono input, you will overload most hi-fi amplifiers and get loud but distorted sound. For similar reasons, to record Mac sounds, connect the audio output port to a tape recorder's auxiliary input rather than to the microphone input.

The sound generators

The sound generator in the small Macs can, in principle, produce frequencies from below audibility to 11 kHz (kilohertz) with four simultaneous voices. In practice, you cannot get more than about 5 kHz effectively. The main CPU runs the sound generator along with everything else it does, so generating complex sounds will slow functions such as screen updating.

Although the generator uses digital techniques, the sound it produces is in no way comparable in quality to digital audio recordings of music. Compact disc players, for example, have 16-bit digital-to-analog conversion, whereas the Macintosh generator has 8-bit conversion. Nevertheless, digitized recordings of real sounds are possible. (Listen, for example, to the Airborne game from Silicon Beach Software, San Diego, CA.)

The Macintosh II has a much more sophisticated sound generator than earlier Macs, with a practical frequency limit of about 7 kHz in the 1987 model. A separate sound processor chip manages it, instead of the CPU, so even fairly complex sounds do not detract significantly from the CPU's other activities. The CPU will be slowed by sound generation only if it must perform major calculations before handing instructions over to the sound chip. The sound generator is capable of stereo sound with a sampling rate of 22.25 kHz and an overall signal-to-noise ratio of about 70 db (decibels)—not as good as compact discs, but good nevertheless. Future sound generators should produce CD-quality output.

The Mac II sound chip can be programmed in several ways to produce sound. It can be programmed like a synthesizer: You specify a frequency (pitch) and a wave envelope (a component of timbre). The Mac can also send signals via MIDI (Musical Instrument Digital Interface) to external MIDI instruments and synthesizers, which produce the actual sounds. Or, like a sampled-sound synthesizer, the Macintosh II can be programmed to take in recorded or live sounds, modify them, and then play back the results.

Speech output

The sound generator in the small Macs is flexible enough to synthesize speech; the one in the Mac II does an even better job. Speech output requires no additional hardware, only speech software. The speech synthesis system operates through the audio port as an output device.

The available speech programs—Apple's Macintalk and Smooth Talker, version 2.0 (First Byte, Long Beach, CA)—operate in two parts. The first part translates ASCII text into phonemes, a phonetic representation of the words. The second part, the synthesizer, takes the phonemes and generates the codes that are turned into sound.

The sound generator can synthesize as fast as it speaks. Speed, pitch, and volume can be regulated independently. To aid pitch and inflection, the synthesizers work on one sentence at a time. In a typical application, punctuation is handled with simple rules: At a comma, the voice will pause slightly with a rising inflection; at a question mark, a longer rise occurs; at a period, a relatively long pause occurs. Decimals are read as "point."

Because English has ill-defined pronunciation rules with many exceptions, more elaborate preprocessing to define inflection and timing might be needed to improve the speech quality. Users will be able to write programs that preprocess text, but they will not be able to modify Macintalk or Smooth Talker directly.

SmoothTalker uses a proprietary synthesis technique and features a male and a female voice. You can see and edit the intermediate phoneme output and call the synthesizer from within a programming language. Twelve hundred rules govern pronunciation; an exception dictionary stores words that do not follow the rules. SmoothTalker is much larger (30 KB + 26 KB for the two parts) than Macintalk but produces better-quality speech.

Macintalk uses formant (resonant components of vowel sounds) synthesis. Its compact code (7 KB + 15 KB) also stores an exception dictionary. Apple has supplied the program to software developers so that they can incorporate it into their own software.

ADDING MEMORY

The main computer circuit board in all current Macintosh models accepts memory chips mounted in plug-in Single Inline Memory Modules (SIMMs). A standard method for packaging memory chips, a SIMM contains eight chips with the same memory capacity. (SIMMs for IBM PC memory boards often have nine chips; the ninth is used for parity checking.) The memory chips come in these standard capacities: 256 kilobits, 1 megabit, and, in the future, 4 or 16 megabits. A SIMM with eight 256-kilobit memory chips can store a total of 256 kilobytes.

In 1988, 256-kilobit memory chips were fairly cheap and readily available. One-megabit chips were available but expensive, and because they cost more than four times as much as 256-kilobit chips, their price per bit was higher. Four-megabit memory chips will not be available at reasonable prices until about 1990; 16-megabit chips might be practical in 1994.

The Mac Plus and Mac SE can accept up to four SIMMs, for a total of 4 MB of memory; the Mac II can accept up to eight SIMMs. (See the Appendix for specific configurations.) Memory can also be installed in NuBus cards on the Mac II; see page 118.

EXPANSION BUSES

The Macintosh SE and Macintosh II both contain an expansion bus with slots that can accept accessory boards. The slots are built into the main computer circuit board (often called the motherboard). The Mac II bus is much more sophisticated and faster than the SE bus; the two are not compatible.

The Mac SE

The SE motherboard has only one expansion slot, with a 96-pin euro-DIN connector. It supplies all bus signals to and from the Motorola 68000 CPU, timing signals, and power for an accessory card. Apple calls this slot arrangement SE-Bus. An SE-Bus card can be up to 4 by 8 inches in size and draw up to 7.5 watts.

The expansion cards that initially might appeal most to SE users are video cards and accelerator cards. Video cards drive big-screen displays—a full page or more. (See Chapter 4.) Accelerator cards install either a faster 68000 or a Motorola 68020 chip into the SE; the latter can put the raw computing power of a Mac II into the smaller SE chassis. Levco (San Diego, CA), General Computer (Cambridge, MA), and Radius (San Jose, CA) all make such accelerator cards. For each one, there is at least one compatible big-screen video board, so you could install both in an SE.

Despite their apparent advantages, neither a video card nor an accelerator card may be a good investment. If you need a large screen or a faster processor, you will probably be better off getting a Mac II to begin with. An SE with an accelerator card will not match the overall performance of a Mac II, which was designed as a faster computer in every way, not merely in processing speed. (An SE with an accelerator board can be faster than a Mac II in some narrowly defined tests; see below.)

An accelerator card costs about half the difference between a Mac SE and a Mac II, and you will be left with a hybrid machine that software developers might not feel obliged to support. The only real gain is the SE's compactness and portability. And if you get a video card and a large monitor, you won't even have much of a size advantage because the monitors are bigger than the SE itself.

Although the SE's expansion slot can be used for adding RAM, using SIMMs that plug into memory sockets on the motherboard is a better way to add memory because the SIMMs leave the expansion slot free for other uses.

Because the SE has only one slot, several companies are building expansion cards that can accept further expansion cards piggyback. A piggyback arrangement can work as long as the entire ensemble of accessory cards does not overload the SE's ventilation system or power supply. Second Wave, Inc. (Austin, TX) is offering a separate external expansion chassis with its own power supply that also contains more slots. But if you really need more than one slot, it's better to get a Mac II. The range of accessory cards made for the SE will be much more limited than for the Mac II.

The Mac II

The Mac II bus is based on NuBus, which was originally developed at the Massachusetts Institute of Technology and first manufactured by Texas Instruments for a workstation computer, the TI Explorer. The NuBus specification provides for two sizes of expansion cards: 14.4 by 11 inches and 4 by 12.875 inches. The larger cards were used by TI in the Explorer and do not fit inside a Mac II. In early 1988, the Mac II was the only commercial product using the small NuBus cards.

On the Mac II, the 96-pin NuBus connector (IEC 603-2) supplies both bus circuits and power. Power is available at three voltages: +12, −12, and +5 volts; the −5.2 volts called for in the original NuBus specification is not available. The total power available to all expansion cards combined is 102 watts.

The fastest data transfer rate possible on the Mac II NuBus is 37.5 MB per second in a short burst transaction. Overall memory access is slowed (in jargon, through wait states) to give the video card time to operate. (Unlike the Mac II's video RAM, the Mac SE video RAM is part of main RAM and requires no delays; this is why an SE with an accelerator board can be slightly faster than a Mac II.)

All NuBus cards contain their own configuration information, including the type of card and manufacturer, in a ROM chip that can be as large as one megabyte. This ROM replaces most jumpers and switches that are commonly used for configuring computer boards. It is possible to design a NuBus card to take over the entire computer, even to the point of displacing the 68020 CPU from control of the bus; in this design, the motherboard is merely another NuBus card.

Second Wave, Inc., has announced an expansion chassis that adds NuBus slots to the SE (distinct from their SE-Bus expansion chassis described above). Again, however, because the SE is not really designed to manage NuBus cards, such a chassis should not appeal to most users. If you want to use NuBus cards, get a Mac II.

NuBus supports a total memory of 4 GB, the same as the 68020 CPU; on the Mac II, 2 GB of memory addressing is available. As a practical matter, memory capacity is limited by physical size. The Mac II motherboard accepts eight SIMMs, each with eight memory chips, for a total of 64 memory chips. A Mac II NuBus card using standard construction techniques can hold 128 memory chips. The maximum number of NuBus cards you could devote to memory chips is five (the sixth slot is taken up by the video card). The total amount of memory available depends on the capacity of the individual memory chips. (See Table 7-1.) Note: The Macintosh operating system supplied by Apple in 1987 could use only 8 MB for main RAM. An error in the first Mac II ROMs prevented the computer from recognizing main memory in NuBus cards, but this problem was fixed by 1988.

TABLE 7-1. MEMORY CONFIGURATIONS POSSIBLE IN A MACINTOSH II

	Chip size			
Placement	*256 kilobits*	*1 megabit*	*4 megabits*	*16 megabits*
Motherboard	2 MB	8 MB	32 MB	128 MB
One NuBus card	4 MB	16 MB	64 MB	256 MB
Five NuBus cards	20 MB	80 MB	320 MB	1280 MB
Packed Mac II[1]	22 MB	88 MB	352 MB	1408 MB

[1] A Mac II containing a motherboard and five NuBus cards filled with memory chips.

8: Printers

A printer is the most straightforward way of getting a job out of a computer. Printers for the Macintosh now come in many varieties, work in diverse ways, and can give you results ranging from rough draft to typeset quality.

PRINTER QUALITY

The most critical distinction among printers is print quality, which depends on several factors:

- Resolution, measured in dots per inch (dpi), or the equivalent, if the printer does not use dots. Table 8-1 gives a scale for rating printer resolution.
- Dot size and shape, an important but rarely discussed factor. For best results, dot size should be approximately the inverse, or reciprocal, of resolution; if resolution is 72 dpi, dot size should be $\frac{1}{72}$ inch. Smaller dots will leave white space between dots; larger dots will overlap. Square dots can fill in areas more completely than round ones; round dots work better for curved edges; oval dots are unsatisfactory. Most printers produce ragged dots; most typesetting machines produce clean ones.

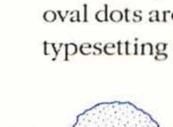

Magnified 30 times, a LaserWriter dot (left) is clearly much smaller than one printed by an ImageWriter.

- Contrast between the printed ink and the paper.
- Uniformity of print quality over the page. Does the printer leave blank areas completely blank? Are black areas even?
- Reflective quality. The printed areas should be matte, that is, free of distracting reflections.

TABLE 8-1. PRINTER RESOLUTION

Rating	*Printer types*	*Comments*
Very low quality: fewer than 125 dpi	Ordinary impact dot-matrix printer; Apple ImageWriters in Fast mode.	Unacceptable for prolonged reading.
Low quality: 125–250 dpi	Most multiple-pass dot-matrix printers; ImageWriters in Best mode; "near-letter-quality" or "correspondence-quality" printers.	Unacceptable for prolonged reading. Cannot do an acceptable tapered serif; sans serif or square serif is better.
Medium quality: 250–500 dpi	A few premium impact dot-matrix and thermal-transfer printers, but their dots overlap; most laser printers and related designs.	With carefully designed fonts, often acceptable for prolonged reading. Cannot do finely tapered serifs well.
High quality: 500–1000 dpi	Formed-character daisy-wheel printers achieve the equivalent of about 700 dpi but lack precise positioning; lower-end typesetting machines.	Suitable for prolonged reading, with some reservations; daisy-wheel printers are marginal. Typesetting machines can reproduce some tapered serifs adequately.
Very high quality: more than 1000 dpi	High-end typesetting equipment only; no microcomputer printer, now or in the near future.	The only choice for long-term reading. Can reproduce all typefaces gracefully, especially at more than 2000 dpi.

There are three broad classes of printers that will work with a Macintosh: conventional dot-matrix, daisy-wheel, and laser. Conventional dot-matrix printers build images out of dots printed on paper. Daisy-wheel printers print already-formed characters one by one. Laser printers, based on xerographic technology, also build images out of dots, but the dots are much finer than those produced by conventional dot-matrix printers.

Many laser printers, including all but one of the LaserWriters from Apple, interpret a special computer language called PostScript, developed by Adobe Systems (Palo Alto, CA), which can handle text as well as sophisticated graphics with great accuracy and flexibility. PostScript requires considerable computing power to process; laser printers have their own microprocessors, memory, and more raw computing power than a small Macintosh.

DOT-MATRIX PRINTERS

The most common microcomputer printers—conventional dot-matrix—produce images by impact: A printing mechanism strikes an inked ribbon or a carbon-coated tape, driving it against the paper and leaving an imprint. Characters are formed from tiny dots. A print head, containing a set of vertically arrayed wires that dart in and out, travels across the paper, pressing the inked ribbon against the paper as it moves. The wires are driven by solenoids, a type of electromagnet.

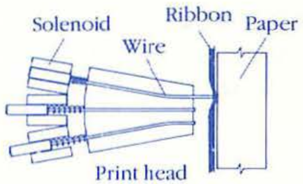

In an impact dot-matrix printer, solenoids drive tiny print wires against ribbon and paper.

Low-cost dot-matrix printers all produce images that look less crisp than the screen display. The printed page lacks the screen's high contrast and uniformity, and the printer's round dots overlap or leave gaps, unlike the pixels on the screen, which are uniform in size and shape and butt together smoothly. In "near-letter-quality" modes, most dot-matrix printers produce dots that are so large that they smudge fine detail. Moreover, the printed image fades as the ribbon wears out.

The ImageWriters

When the Macintosh first appeared, Apple also introduced the ImageWriter printer to go with it. The original design of this basic dot-matrix printer has since been updated to produce the ImageWriter II; much in the descriptions that follow, however, also applies to the original ImageWriter. In addition, Apple has come out with another model, the ImageWriter LQ, whose performance and print quality lie part way between the ImageWriter and ImageWriter II and the LaserWriters.

The ImageWriter II

Today's standard low-cost dot-matrix printer for the Macintosh is Apple's ImageWriter II. All ordinary programs will drive it (and also the original ImageWriter), and many accessories are designed specifically to work with it. It can print all Macintosh screen fonts and graphics as well as replicas of the screen called screen dumps. The ImageWriter II prints in three modes: Draft, Standard Quality (called Faster in printing menus), and Near Letter Quality (Best). Which modes are available depends on particular application software; many programs offer all three. The ImageWriter print head contains nine wires; dot size is 16 mils ($^{16}/_{1000}$ inch, or 0.4 mm).

For conventional dot-matrix printing, including all modes except Draft, Macintosh software first takes data (such as font information) stored in the System file and creates a bit map in RAM of the material to be printed; it then sends the bit map to the printer. The bit map is created in bands, which are printed by successive passes of the print head; the printer pauses while the software creates the next band. In Best mode on the ImageWriter II, for example, each band corresponds to four passes of the print head; a page contains 47 bands.

Draft is the fastest mode on an ImageWriter; its quality is also the lowest. Use Draft when you want quick results and don't care how the page looks. In Draft, the Macintosh does not create a bit map of the screen, but sends only a code for each letter and number; the printer uses its own built-in font. For this reason, you can print only text and numbers in draft mode, not graphics. To get even word spacing on the printed page,

©Martin Easley

use a monospaced font, such as Monaco, on the screen. (In a monospaced font, all the letters take up the same amount of horizontal space, as on manual typewriters.)

The ImageWriter II printing in draft mode.

Faster and Best modes can print with proportional spacing, in which the width of each letter varies so that, for example, a capital *M* takes up more space than a capital *I*, and a lowercase *w* more space than a lowercase *l*; numerals are all the same width so that columns of numbers will line up properly.

Faster mode prints at 72 dpi vertically and 80 dpi horizontally, which is similar to the Mac screen (72 dpi both horizontally and vertically). Nearly all software will print in this mode without complications. Unlike in Draft mode, Macintosh software printing in Faster mode sends a bit map of the screen image to the printer, giving you pixel-for-pixel correspondence with the screen except for the 11 percent horizontal distortion. You can get rid of this distortion in many programs by choosing Tall Adjusted printing. But Tall Adjusted printing moves the paper more slowly and forces a compromise in print head positioning; the results can be a little untidy.

In Faster mode, both the screen display and the printed page are generated from identical font information stored in the System file. The dots do not overlap, leaving gaps and making the result very low quality.

The ImageWriter II printing in faster mode.

At 144-by-160 dpi, Best printing offers twice the linear resolution of the Mac screen and of Faster printing. To achieve this, the printer makes two passes over the paper, printing twice as many dots horizontally as in Faster mode. On the second pass, the paper is rolled up the height of half a dot, so the dots overlap on the page by 50 percent; quality is low nonetheless.

If you print many pages with large areas of black in Best mode, the solenoids in the print head can overheat. Print such pages one at a time, and let the print head cool down in between.

The ImageWriter II printing in best mode.

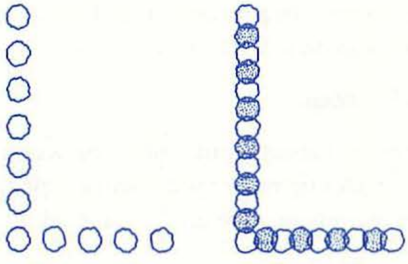

Best mode prints twice as many dots per inch as Faster mode.

Because of limitations in screen resolution, you cannot see on the screen exactly what Best mode printing will look like; the extra dots are added during printing on paper. Each printed character contains four times as many pixels as the original screen character (twice the size vertically and horizontally means four times the area). To create the extra pixels, printer driver software (also called the printing resource), which is part of Mac system software, uses font information in the System file for characters that are twice as large as the ones you are using on the screen. The bit map for printing a 12-point character is thus actually created from a bit map of a 24-point character. If the appropriate double-sized font has not been installed in the System file, the printing resource looks for a font four times as large. If that doesn't exist, the print head shifts slightly and simply double-strikes the paper.

Although an ImageWriter can print the screen fonts—such as Times and Helvetica—that correspond to a LaserWriter's PostScript fonts (see page 142), the spacing and positioning will not match a LaserWriter's output. Thus an ImageWriter is not satisfactory for proofing pages that will be printed on a PostScript printer such as a LaserWriter. Furthermore, ordinary screen fonts—such as New York and Geneva—produce better results on an ImageWriter.

An ImageWriter can print in color with a color ribbon containing four inks: black, yellow, magenta, and cyan (blue-green). It prints color like any printing press, by a process called subtractive mixing, in which light is absorbed, or subtracted, by each different ink color. The human eye sees light that passes through the ink and reflects off the white paper beneath. Unfortunately, this mixing usually produces ugly results on the ImageWriter; color printing is best restricted to objects that do not physically overlap, such as isolated bands of color. In any event, color ribbons are expensive.

The ImageWriter LQ

The new ImageWriter LQ is faster—by nearly two times—and more expensive than the ImageWriter II. Its print head contains 27 print wires, compared with only nine wires on the ImageWriter and the ImageWriter II. Like the ImageWriter and ImageWriter II, the ImageWriter LQ has three modes: Best, Faster, and Draft. In its Best mode, it offers 216-dpi vertical and horizontal resolution. The printer is capable of printing 320 dpi horizontally, although most Mac software will only drive it at 216 dpi.

For the ImageWriter LQ to print in Best mode, you need to install in the System file a font that is three times the printed size (to print a 12-point font, you need to install the 36-point size of that typeface; to print a 24-point font, you need the 72-point size). Because such large fonts take up considerable storage space, the ImageWriter LQ functionally requires a hard disk drive. Fonts for the ImageWriter and ImageWriter II work with the LQ, but few of these come in the large sizes necessary for Best quality printing. If the LQ proves popular, font vendors may begin developing large sizes of their fonts.

Apple supplies Times, Helvetica, Courier, and Symbol in 9-, 10-, 12-, 14-, 18-, and 24-point sizes to print with Best quality. These fonts look similar to the corresponding fonts

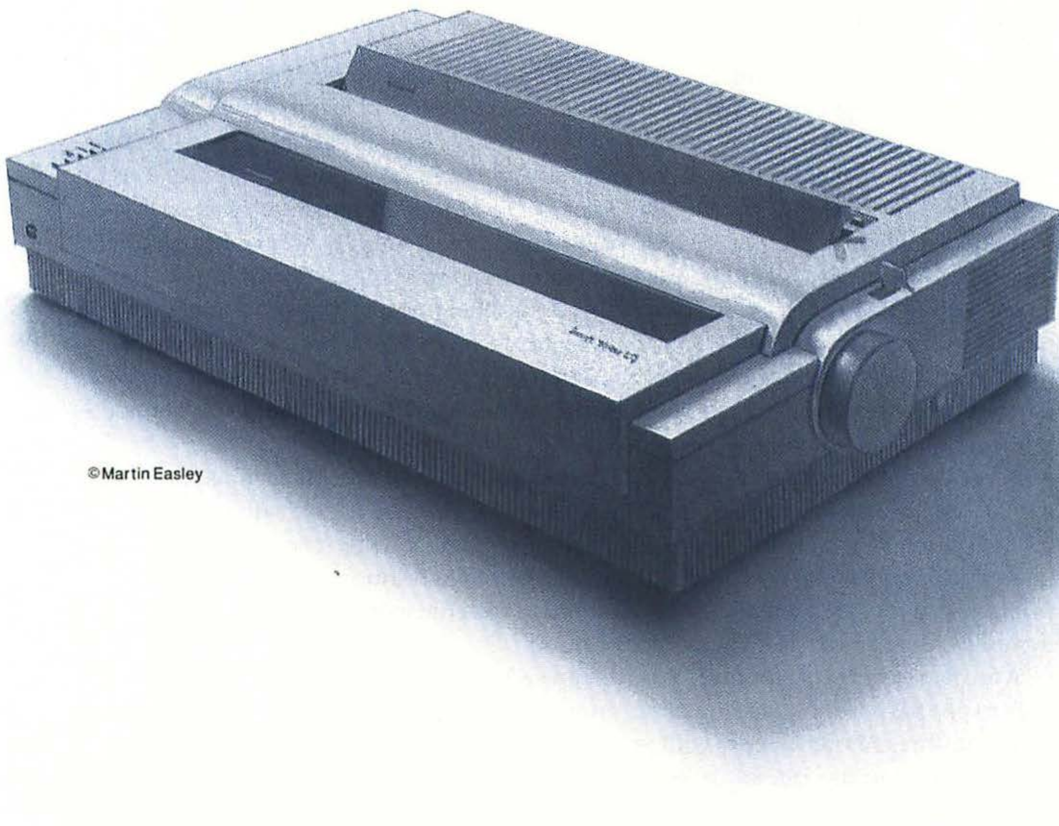

©Martin Easley

supplied with the LaserWriters, but they do not print with identical spacing. The LQ, like the ImageWriter and ImageWriter II, cannot be used for precise proofing of a page that will go to a PostScript printer or typesetter for final output.

This sample text was printed by the ImageWriter LQ in Draft mode.

This sample text was printed by the ImageWriter LQ in Faster mode.

This sample text was printed by the ImageWriter LQ in Best mode.

Apple offers several accessories for the ImageWriter LQ. A sheet-feeder attachment can accommodate up to three bins, so you can print on letterhead, continuation sheets, and plain paper without attending to the printer. You can also get a color ribbon for printing color as the ImageWriter II does. And an AppleTalk option lets several people share a single LQ.

Other Dot-Matrix Printers

Many alternatives to the ImageWriter, ImageWriter II, and ImageWriter LQ are now available, including clones that behave like ImageWriters and use ImageWriter printer driver software. You might want to use a printer other than the ImageWriters because you already have one, because of compatibility with other computers, or for reasons of cost. There are some disadvantages to doing so, however: Macintosh software was specifically designed for an ImageWriter, so other printers generally do not work as well; they print different dot sizes and vary in their minimum spacing, so your finished copy looks distorted. Within a given price range, most other impact dot-matrix printers work more slowly than the ImageWriters, and you may not be able to print in all modes with them. A parallel printer (which needs a Centronics parallel port) will need a serial-to-parallel converter. And only the ImageWriter and ImageWriter II accept the ThunderScan scanner. (See Chapter 17.)

Several low-cost clones of the ImageWriter and ImageWriter II, such as the Epson AP-80, are now on the market. In addition, some companies have made retrofit ROMs for their printers, which are supposed to make them ImageWriter-compatible. Most dot-matrix printers are not, however, and require their own printing resource. To use these printers, you must replace the ImageWriter icon with the icon of the new printer in the System Folder of each disk you use for printing.

A few application programs bypass the ImageWriter printing resource altogether and thus might have trouble driving other printers. These programs themselves create the codes an ImageWriter needs, without requiring the ImageWriter printing resource. If you are trying to use one of these programs with a non-ImageWriter printer, you will probably have problems, but neither the software company nor the printer company may be able to help because, in most cases, neither company will be familiar with the other's product.

For non-ImageWriter printers, the SoftStyle subsidiary of Phoenix Technologies (Honolulu, HI) offers the most comprehensive set of printer drivers. Its program, Printworks for the Mac, supports more than 20 printers, including the ImageWriter, and comes with a printer buffer (see page 134) and a screen preview feature, which shows you where pages break and takes the guesswork out of printing only certain pages. If your application programs, like most, don't have page previewing, this feature makes Printworks valuable for an ImageWriter as well.

Printing precision among all dot-matrix printers varies considerably. Middle-of-the-road printers can sometimes work better than either expensive high-speed printing mechanisms or cheaply made printers. The most common problem is in printing vertical lines. Any printer can put down a straight horizontal line, but many of them have trouble matching up segments of a vertical line on repeated passes of the print head. Do not buy a printer for the Mac unless you have seen it working with a Mac and are satisfied with the results.

LETTER-QUALITY PRINTERS

For all ImageWriter models, Near Letter Quality, or Best, mode falls far short of the letter-quality output of a good electric typewriter. Printers that can give true letter quality use daisy wheels (or their close cousins, thimbles) with molded embossed characters. A print hammer strikes one of the petals on the daisy wheel, pressing it against the ribbon and paper and printing an entire character at once. In the most sophisticated daisy-wheel printers, the hammer strikes the letter *M* harder than it strikes a period, making the overall impressions even. As with typewriters, print sharpness and resolution depend on the precision of each molded character. Since daisy-wheel printers generally use carbon ribbons, which are discarded after one use, ribbon condition matters less than it does for dot-matrix printers.

Daisy wheel print DAISY WHEEL PRINT

This is a sample of 10 pitch type.

ABCDEFGHIJKLMNOPQRSTUVWXYZ

Although daisy-wheel printers produce clean type, they have many disadvantages. They are slow and noisy, typically printing 12 to 40 characters per second, although many are faster than the ImageWriter or ImageWriter II in Best mode and lack the high-pitched whine of a dot-matrix unit. They are limited to one typeface at a time; to change typefaces you must change the daisy wheel, which slows the printing process considerably. Type size is limited by the size of the daisy wheel itself, so you can only get the print sizes available on normal typewriters.

In addition, Macintosh software may not display on screen the correct line lengths for a particular daisy-wheel printer, making it impossible for word processors to hyphenate correctly at the ends of lines. (One word processor that does show true line lengths for daisy-wheel printers is Microsoft Word.) Most important, daisy-wheel printers are only good for characters; they cannot print graphics effectively. Some daisy-wheel printers use the period to build up graphics images, in effect imitating a dot-matrix printer, but the process is slow, and the period petal on the daisy wheel wears out quickly.

Several companies have produced software and hardware for connecting popular daisy-wheel printer models to a Macintosh. If you want to buy a daisy-wheel printer, look for a model compatible with the Diablo 630 printer. In a pinch you can use an electronic typewriter with suitable interface hardware and driver software. But electronic typewriters are very slow, and nearly all of them require manual paper feeding, one sheet at a time.

LIVING WITH CONVENTIONAL PRINTERS

Printing with conventional, rather than laser, printers involves a variety of considerations beyond putting paper in and watching finished copy come out:

- Paper handling and switching. The paper-moving mechanism, or tractor feed, can be hard to thread on many printers, making it difficult to alternate between formfeed paper and single sheets or envelopes. Try before you buy.
- Speed. Manufacturers usually quote the most optimistic measurement—the highest burst speed of printing a single line. With time included to advance the paper to the next line or page, however, overall printing speed is typically only a third to a half of the claimed speed. Many low-cost printers do not have true formfeed and can advance only a line at a time whether printing or not; better printers can move an entire page quickly when not printing.
- Sheet feeders. Formfeed paper is reliable but inconvenient; you have to separate the sheets when you are done, and you cannot use letterhead or other special paper unless it comes in a formfeed version. A sheet feeder inserts ordinary cut paper into a printer automatically. The sheet feeders for the ImageWriter LQ are better than average, but even the best sheet feeders can perform poorly because of variations in paper surfaces and dirty rubber rollers; jams and misfeeds are common. If you want to be able to walk away from the printer when you are printing a long document, better stick with formfeed paper.
- Double-sided printing. Very few printers made for microcomputers can print on both sides of a sheet automatically. If you need double-sided printing, sheet-feeding is easier to work with than formfeeding. Print your document once, then shuffle the paper so that the first page 1 will be backed by the second page 2, and so on. In other words, for a six-page document, stack the

pages in the following order: 2, 1, 4, 3, 6, 5. Then print again on the back side. Be sure you know in which order your software instructs the printer to print the pages (whether it prints the first or the last page first). This trick will produce a minimum of two copies of your document. A simpler but slower way is to feed paper to your printer manually, pausing between printing each page to turn the paper over.

- Wide-carriage printers. Many printers, including the ImageWriter II, are available with both standard (10-inch, or 25-cm) and wide (14–15-inch, or 35–38-cm) carriages. Printers with wider carriages will let you print larger images, but your software has to be able to take advantage of the extra width. MacWrite version 4.6, for example, does not let you create documents wider than a small Mac screen.

Ongoing Costs

The total cost of a printer includes not only its purchase price but the cost of keeping it fed with ribbons and paper. Although the price of paper is the same for any printer, ribbon costs vary widely. You might also need to replace the print head from time to time.

Printer ribbons

When considering a printer, find out whether it uses standard ribbons available from many sources or the manufacturer's own, usually overpriced, ribbons. In either case, carbon ribbons are expensive because they can be used only once.

Dot-matrix printers generally have cloth (nylon) ribbons. Unlike a carbon ribbon, which is fixed in length, a cloth ribbon forms an endless loop inside its case. How long the ribbon lasts thus depends on your tolerance for slowly fading ink. Many users keep an old ribbon for routine printing and a nearly new ribbon for important jobs. If you need the cleanest possible output from an ImageWriter, look for a carbon ribbon.

Depending on printer and ribbon design, you have several choices for replacing cloth ribbons:

- Buy a new one.
- Replace only the ribbon, if the case can be opened and you don't mind messy fingers. (In some printers, you can simply turn the ribbon over and print on another portion of it.)
- Reink the ribbon using a kit advertised in computer hobby magazines. You should probably reink only a few times; a frayed ribbon can damage print heads by snagging and bending the print head wires.
- If you are truly desperate, one simple way to rejuvenate a fading cloth ribbon is to open the case and spray a little WD-40 or similar lubricant on the ribbon. Let it sit for a few minutes after spraying. The lubricant acts as a solvent, spreading more of the ink over the ribbon's surface.

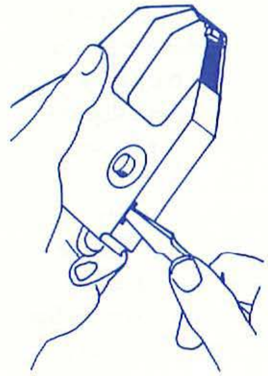

Open the ribbon case with a small screwdriver.

One last hint: Don't stockpile ribbons; because the ink dries out, keep enough for only a few months rather than a few years.

Paper

Formfeed computer paper has extra margins with holes that fit into a printer's tractor feed; it is available in many styles and weights. For normal use, buy 20-pound, 9½-by-11-inch paper; tearing off the perforations will leave 8½-by-11-inch paper. You can also set Mac software for the slightly narrower and longer A4 paper size common in Europe and Japan. The ImageWriter driver also offers a wide printing format that prints sideways, with the top line of your document running down the right side of the paper. (The paper is loaded normally.) Table 8-2 lists paper types and sizes.

TABLE 8-2. PAPER SIZES

B4 (international)	10.1 by 14.3 in.	(257 by 364 mm)
Legal	8.5 by 14 in.	(216 by 256 mm)
Foolscap	8.5 by 13 in.	(216 by 330 mm)
Folio	8.3 by 13 in.	(210 by 330 mm)
A4	8.3 by 11.7 in.	(210 by 297 mm)
Letter	8.5 by 11 in.	(216 by 279 mm)
B5	7.2 by 10.1 in.	(182 by 257 mm)
Ledger	11.0 by 17 in.	(279 by 432 mm)
International fanfold	8.3 by 12 in.	(210 by 304 mm)

Removing the perforated margins of ordinary formfeed paper leaves the edges slightly ragged, but you can buy more expensive paper with much smaller perforations and much cleaner edges, suitable even for letterhead.

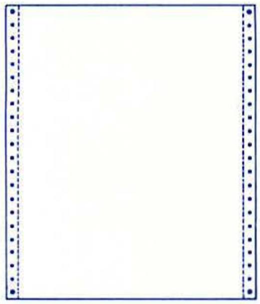

Formfeed paper.

Mailing labels also come on formfeed stock. A word of warning: Never try to roll formfeed labels backward through a printer, or they will jam. The labels are normally sized for six lines per inch; most word processors can be set to print at this spacing. The best program for printing on labels is Silicon Press from Silicon Beach Software (San Diego, CA). It enables you to specify label size, print a single label many times, or print a mailing list in which every label is different. (Silicon Press also works with a LaserWriter and other PostScript printers.)

Envelopes are a nuisance because you must alternate them with ordinary sheets and therefore attend to the printer frequently. You can copy addresses to a separate file for printing in a batch, but this is also clumsy. Window envelopes are an excellent solution because you can simply use the printed address in the letter itself.

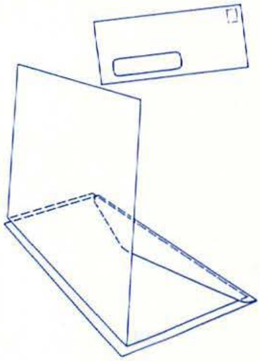

To fold a letter for a window envelope, use the envelope flap as a guide.

To get copies of your printed pages, you have several choices:

- Make a photocopy. This saves the trouble of changing paper or inserting carbons.
- Use carbon paper with single sheets. Impact printers can usually handle at least one carbon copy. If ordinary carbon paper does not make dark enough copies, try Carter's Super Midnight copy film.

- Use multipart forms paper, which includes a carbon copy or equivalent. With multipart or carbon paper you may need to increase the print intensity, and this will wear down the print head a little faster.
- Print a second copy. Although quick with laser printers, this is time-consuming with conventional printers. On an ImageWriter, you could save time by selecting draft mode for the second copy.

Printer Buffers

Computers can feed information to conventional printers much faster than the printers can print, which means that you spend a lot of time waiting for the printer to finish a job. Printer buffer software, also called printer spooler software, takes the printing instructions from the computer at high speed, stores them in RAM or on a disk, and then feeds the information slowly to the printer as required. Once everything goes to the buffer, the computer "thinks" the printer is done and becomes available for other work.

Buffers can store printing instructions in several different places:

- In main RAM or on a disk in the computer
- In the RAM of a freestanding buffer box that goes between the computer and printer
- In RAM or on a disk in the printer
- On a disk in a network

For a single-user system, a printer buffer in the computer's main memory or on a disk is usually the best choice. On a network, printer buffering is normally done by the main file server.

Printer buffers are handy however they work, but to be effective they must have one key feature: You must be able to erase their contents—in case you start printing and realize that you have made a formatting error, for example. If you cannot erase the buffer, the only way to avoid printing your entire document, errors and all, is to turn off the computer and printer. (Typing Command-period does not affect buffers.)

The fanciest and most expensive buffers are the freestanding boxes; some include an erase button as well as a reprint button, so if you want another copy, the buffer sends the information to the printer again. But it rarely makes sense to buy a buffer box; instead, add more memory to your computer, which you can then use for many other tasks as well.

Printer Stands

A printer stand can make your life easier by keeping formfeed paper from becoming tangled on the floor. If you have room, build or buy a stand with space below the printer for an entire box of paper. If not, you can find a variety of simple stands that will fit on a tabletop. If you feed paper from a box and cannot put the box underneath the printer (because the printer is on a filing cabinet, for example), put a small board across the paper box to receive the output. In all cases, check that paper leaving the printer doesn't interfere with paper feeding in.

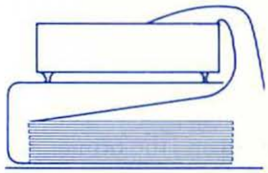

A printer stand makes it easy to use a printer on a table.

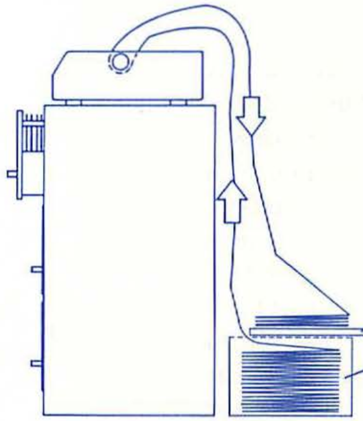

Printer Sound Hoods

Any impact printer—dot-matrix or daisy-wheel—makes too much noise for anyone to concentrate nearby; you can't even carry on a telephone conversation in the same room. A sound hood—a large, bulky box with a sound-absorbing lining—can make life with a printer much more peaceful. If you use one, be sure that the printer still has adequate ventilation to dissipate heat. Check paper handling when the hood is on and closed; sound hoods often make it difficult to get at the paper.

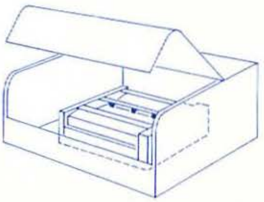

Sound hood.

You could also banish a noisy printer to another room or to a closet. Or put the printer into a cubicle lined with sound-absorbent material such as a cloth wall hanging or curtain. The best, but most expensive, material is Sonex acoustic foam, used in sound-recording studios (made by Illbruck, Minneapolis, MN).

Sharing a Conventional Printer

If you need to connect several Macintoshes to a single conventional printer, you can install switching boxes between the printer port of each Macintosh and the printer. A better solution is to plug an optional AppleTalk card into an ImageWriter II or ImageWriter LQ and run it as a network printer.

THE LASERWRITERS

Laser printers print by means of xerographic technology and thus differ fundamentally from impact printers. Rather than stamping or inking images onto paper, they use electrostatic charges, revolving drums and cylinders, powdery black toner, heat, and pressure to create a photocopylike image. The amazing range of processes built into laser printers makes them the most complex technology in offices today.

For image stability and geometric accuracy, a laser printer must print an entire page once it starts. Unlike ordinary printers, it cannot pause in midpage to receive more information; all the information needed to print a page must be available at the beginning. No standard microcomputer communication link can move all the printing information fast enough for the computer to drive the laser beam directly, so any laser printer needs its own processor and memory—RAM, ROM, or both. The need for a processor and memory is even greater for laser printers that interpret the PostScript page-description language. (See Table 8-3 for a comparison of some laser printers.)

So far, Apple has produced five LaserWriters:

- The LaserWriter (introduced in 1985): based on Canon's LBP-CX printing mechanism and containing its own PostScript processor.
- The LaserWriter Plus (introduced in 1986): essentially the same as the LaserWriter but with a larger ROM for additional fonts (see below).
- The LaserWriter II series (three models introduced in 1988): based on Canon's LBP-SX printing mechanism. The LaserWriter IINT and IINTX interpret PostScript. The simplest and cheapest model, the LaserWriter IISC, works with QuickDraw instead of PostScript.

All the PostScript-based LaserWriters work with the Macintosh mainly through the AppleTalk local area network, and so can be easily shared. They also work with other computers, either using AppleTalk or direct connection to an RS-232C port; through that port, a LaserWriter can operate either as a PostScript printer or as a conventional text printer, accepting the most common Diablo 630 printer commands. The LaserWriter IINTX also accepts Hewlett-Packard LaserJet Plus commands.

TABLE 8-3. COMPARISON OF SOME LASER PRINTERS FOR THE MACINTOSH

	CPU	Speed	RAM	ROM	Built-in font families	PostScript?
LaserWriter	68000	12 MHz	1.5 MB	512 KB	4	Yes
LaserWriter Plus	68000	12 MHz	1.5 MB	512 KB	10	Yes
LaserWriter IISC	68000	7.45 MHz	1 MB	16 KB	4	No
LaserWriter IINT	68000	12 MHz	2 MB	1 MB	10	Yes
LaserWriter IINTX	68020	16 MHz	2–12 MB	1 MB	10	Yes
QMS-PS 810	68000	16 MHz	2–3 MB	1 MB	10	Yes

How the LaserWriters Work

For PostScript laser printers:

- The printing process begins when your application software sends information that describes the page to be printed to a utility program, the laser printer printing resource, or driver, in the Mac's System Folder. The printing resource translates this information into PostScript. (Some applications create PostScript commands directly.)
- The printing resource then sends the PostScript commands to the printer via the AppleTalk network.
- The PostScript processor interprets the PostScript commands, converting them into a dot-matrix image with 300 dots per inch. The image is laid out in a series of lines, much like a television picture, which modulates a laser beam.

Meanwhile, the main printing process plays out on a rotating photoconductor drum inside the printer (PostScript and non-PostScript printers):

- A preconditioning lamp illuminates the drum's surface, neutralizing any charges left over from printing the previous page.
- The charging corona, a fine wire held at a high voltage, ionizes and lays down air molecules as an even layer of negative electric charges on the photoconductor surface.
- A rapidly spinning mirror deflects the modulated laser beam across the photoconductor. A compensator lens corrects the geometry of the laser beam as it moves across the drum.
- Wherever the beam strikes, the photoconductor, as its name implies, conducts electricity so that the negative charges drain off what will become the image's black areas, leaving these areas at a positive potential relative to the rest of the

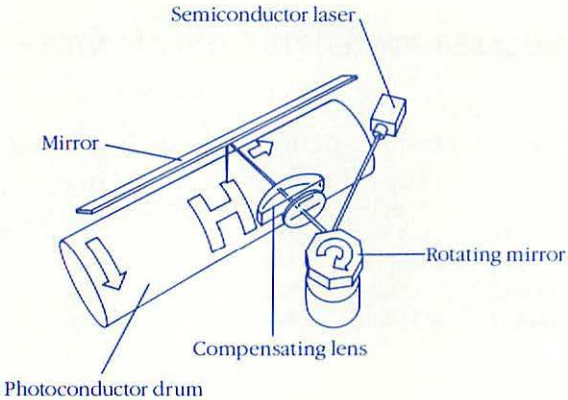

image. This creates an electrostatic latent image, much like the latent image on photographic film after it has been exposed.

- The toner (a monocomponent toner) consists of magnetic particles covered with meltable plastic resin and carbon black. An organic polymer coating on the toner particle creates negative charges by friction (triboelectricity) against other toner particles. The toner is fed into a small gap between a magnetic blade and the developing cylinder.
- The rotating developing cylinder contains a fixed magnet; the concentrated magnetic field creates a toner "curtain" between the blade and magnet. As it turns, the developing cylinder picks up from the curtain an even coating of toner, which sticks to the cylinder because of the fixed magnet.
- The developing cylinder and the photoconductor drum rotate together very closely but don't quite touch. A rapidly alternating electric field between the cylinder and drum creates a tiny cloud of toner along this gap. Areas on the drum that are negatively charged repel the toner, pushing it back to the cylinder, where it is held by the magnet. Discharged areas are positive with respect to the toner and attract it, so the toner jumps the gap and sticks to the photoconductor, developing the latent image.
- A transfer corona charges the paper positively, so that when the paper rolls against the photoconductor drum, the negatively charged toner sticks to the paper.
- The toned image is fused to the paper by a combination of heat and pressure.
- A cleaner blade wipes excess toner off the drum, and the process begins all over again.

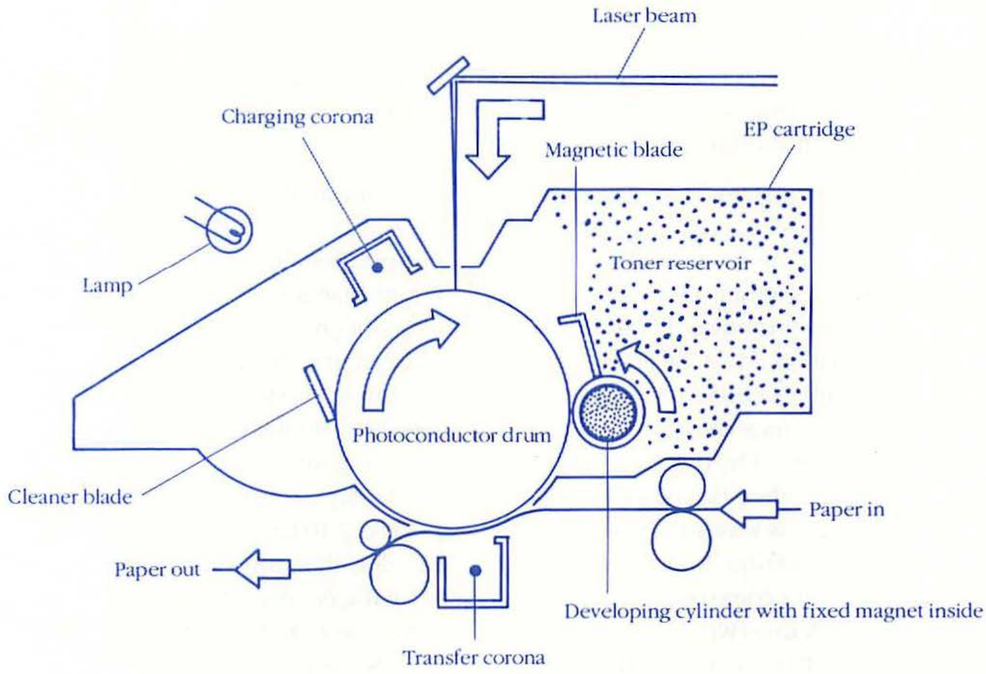

The Page-Image Buffer

At 300 dots per inch, a page contains nearly 8,000,000 dots. A printer capable of full graphics must therefore have 8,000,000 bits, or about a megabyte, of memory as a page-image buffer.

A LaserWriter with PostScript can print an 87-square-inch area laid out in any of the following formats:

Page type	*Paper size (inches)*	*Image size (inches)*	*Pixels*	*Buffer size (bytes)*
Letter	8.5 by 11	8 by 10.92	2400 by 3276	982,800
Lettersmall	8.5 by 11	7.68 by 10.16	2304 by 3048	877,824
Legal	8.5 by 14	6.72 by 13	2016 by 3900	982,800
A4	8.27 by 11.69	7.79 by 11.08	2337 by 3324	971,000
A4small	8.27 by 11.69	7.47 by 10.85	2241 by 3255	911,800
B5	6.93 by 9.84	6.45 by 9.76	1935 by 2928	708,210

The printable area is always centered on the page. The LaserWriter printing resource normally selects a printable page size that depends on which size paper tray has been installed. If a letter or A4 tray is installed, the printable area will be Lettersmall or A4small; if a legal or B5 tray is installed, the printable area will be Legal or B5.

How PostScript Works

All LaserWriters with PostScript, and most other PostScript laser printers, work according to the following outline. Depending on what you are printing, PostScript provides the printer with several kinds of information:

- Font information, including style, size, and printing angle. Size and angle are continuously variable.
- Character information. Once a font has been specified, the printing resource sends only simple character codes until the font changes.
- Graphics primitives. For drawing a line, PostScript gives the printer information about the line's attributes, such as length, width, and orientation; the printer's PostScript processor then figures out where to put the pixels to print the image at 300 dpi. Describing an object's attributes rather than building it dot by dot saves considerable processing time.
- A bit map, the pixel-by-pixel image. ImageWriter-style printing at 72 dpi can be done this way; each pixel is treated as a block of 16 LaserWriter dots. Because 300 dpi is not a simple multiple of 72 dpi, bit-mapped images on a LaserWriter come out four percent smaller than the equivalent ImageWriter output. A LaserWriter can also print at the same size as an ImageWriter, but the image will be a little distorted because the PostScript processor must interpolate bits to make the size come out correctly. With appropriate software, you can also print bit maps at the full 300 dpi, but you should do this only if absolutely necessary because it takes so long to send information with so much detail.
- Gray-scale information. Macintosh software can send information about grays along with the bit map, and the PostScript processor will create a halftone pattern to simulate the grays. (See Chapter 12.)

Four basic typeface families are built into the ROM of nearly all PostScript printers:

- Times Roman, Times Italic, Times Bold, Times Bold Italic
- Courier, Courier Oblique, Courier Bold, Courier Bold Oblique
- Helvetica, Helvetica Oblique, Helvetica Bold, Helvetica Bold Oblique
- Symbol

More fonts have been added to the ROM in the LaserWriter Plus, the LaserWriter IINT, the LaserWriter IINTX, and other enhanced PostScript printers:

- Palatino, Palatino Italic, Palatino Bold, Palatino Bold Italic
- ITC Avant Garde Gothic Book, ITC Avant Garde Gothic Book Oblique, ITC Avant Garde Gothic Demi, ITC Avant Garde Gothic Demi Oblique
- ITC Bookman Light, ITC Bookman Light Italic, ITC Bookman Demi, ITC Bookman Demi Oblique

- Helvetica Narrow, Helvetica Narrow Oblique, Helvetica Narrow Bold, Helvetica Narrow Bold Oblique
- ITC Zapf Chancery Medium Italic
- ITC Zapf Dingbats
- New Century Schoolbook, New Century Schoolbook Italic, New Century Schoolbook Bold Italic

Of this list, Helvetica Narrow is not a distinct typeface; it is generated by a scaling algorithm that squeezes the standard Helvetica font. ITC stands for International Typeface Corporation, a major source of fonts in the typesetting world. *Demi* means demibold, intermediate between normal and bold. Both oblique and italic fonts are slanted, but oblique is essentially generated by tilting letters, whereas italics are actually a different font design. For more information about fonts, see Chapter 12.

How PostScript stores and finds fonts

PostScript's basic font information consists of outlines rather than bit maps. One set of outlines is used to generate all sizes, from too-small-to-read to larger-than-the-paper. PostScript itself sets no limits, although Macintosh applications are more restrictive. A few programs can print only font sizes that correspond to a screen font installed on your computer. Many programs offer a fixed range of font sizes; the most flexible programs let you select any size, ranging from 4 point, which is the smallest legible size on a LaserWriter, to 127 point, the largest screen size supported by the Macintosh ROM. A few programs let you select fractional point sizes (9.5 point, for example) or fonts larger than 127 point. Besides the font outlines stored in the LaserWriter ROM, new font outlines (typically requiring 15 to 35 KB per font) can be stored on disk and transmitted to the printer (downloaded) as needed.

Three fonts are stored as bit maps in 36 KB of the LaserWriter ROM, so they are always available: A full ASCII character set in 10-point Courier type, and alphanumerics (letters and numbers) and common punctuation in 12-point Times Roman and Helvetica. When you turn on the power to the printer, a LaserWriter builds up several other fonts by reading the outlines stored in its ROM, creating a bit map for each character, and moving the bit map to the area of its RAM called the character cache. These fonts are 10-point Times Roman and Helvetica, plus lowercase letters for Times Bold and Helvetica Bold in 10-, 12-, and 14-point sizes.

To save memory, the bit map of each character is trimmed to fit within the smallest rectangular block that can contain it. A bit map for a complete 10-point font takes about 10 KB of memory; the storage required goes up with the square of the point size.

For a lettersmall-sized page-image buffer (see table on page 139), the original LaserWriter's 1.5 MB of RAM is partitioned as follows: 878 KB for the page buffer, 150 KB for the character cache, 345 KB for a work area, and the remainder for PostScript's internal processing. The LaserWriter IINT and LaserWriter IINTX have 2 or more MB of RAM.

The additional RAM is mainly used to increase the printer's work area. Page-buffer size is controlled by software; the smaller the buffer, the more memory left for font information and processing. For this reason, the standard printed text area in most Mac programs is smaller than the maximum printable area. If the printer's RAM is expanded, as the RAM in the LaserWriter IINTX can be, the standard text area can also be increased. (Apple does not sell any memory expansion for the original LaserWriters, but independent companies do.)

When you print a character, the PostScript processor checks the character cache; if the desired character is there, the processor sends it to the page buffer. If what you are printing calls for fonts or sizes not in the character cache, the processor creates bit maps character by character, as needed, and puts them in the cache. Individual bit maps are created for all characters that are rotated from the horizontal—for example, if you were printing a special display page. If the character cache fills up, the processor discards the font with the least recently used characters. In most PostScript printers, the character cache can store about 700 12-point characters; a typical page uses about 200 distinct characters.

Characters larger than a certain threshold size are not put into the character cache but are instead created when needed and immediately moved to the page buffer. This conserves cache memory as well as processor effort, because such characters are seldom called for. The size threshold is set by a Postscript command and depends on the memory available in your printer. For the LaserWriter and the LaserWriter Plus, this threshold size is 1200 bytes, which corresponds to about one 24-point character.

Altogether, building up a font in the original LaserWriter takes about $^1/_3$ second for a 12-point character. Thus, the first page in a new font will take longer to print than the following pages. Simply moving a character's bit map from the cache to the page buffer takes only 500 microseconds.

Screen Fonts and the LaserWriters

On the Macintosh, PostScript is primarily a language for describing the printed page, not what you see on the screen. Most Macintosh applications use another language, QuickDraw, to create screen images and then convert QuickDraw to PostScript just before printing. For on-screen display, QuickDraw therefore needs a 72-dpi font that matches the PostScript outlines used by the printer. Preferably, the outline and corresponding screen fonts should look the same, but this is difficult to achieve because so few dots are available for the screen fonts, particularly for sizes below 12 points. (See Chapter 12.)

If the font you are using on screen does not have a corresponding PostScript font in your LaserWriter, the LaserWriter printing resource will either print your text as a bit map (which could take a long time) or, for three common screen fonts, substitute outline fonts: New York prints as Times, Geneva as Helvetica, and Monaco as Courier.

This substitution is not very satisfactory, however, because although each character is printed from PostScript's outlines, the spacing between letters, words, and lines comes from the original screen font and is disproportionately loose.

It is possible to print screen fonts designed for an ImageWriter (for example, Chicago, Venice, and so on) on a LaserWriter, and the LaserWriter's output can look a little better than the ImageWriter's because the dots are square and clean. If you click the box for smoothing in the Print dialog box, some fonts improve. Other, more intricate fonts do not fare so well because smoothing does only 45- and 90-degree straight-line interpolation. Because this interpolation is always done by filling in with black half-pixels, smoothing can give less clean results when printing white characters on black than when printing black on white.

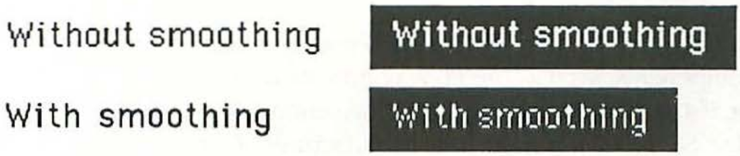

You can also print screen fonts designed for a Postscript LaserWriter on an ImageWriter, but the results are usually unsatisfactory.

The Printer and the Page

At 300 dots per inch, the LaserWriters are, strictly speaking, medium-resolution printers. The quality of the printed page is not as good as with typesetting; it looks too much like a photocopy, and characters lack the crispness of those produced by formed-character impact printers.

Although the laser writes dots of fixed size (4 mils) on the photoconductor drum, the size of the printed pixels varies because the paper surface and the toning and fusing steps all distort the dots into irregular splotches, when they are viewed under magnification on the finished page. Because electrostatic images are unstable (the charged toner particles repel each other on the photoconductor drum), black areas are printed unevenly, particularly in the LaserWriter and LaserWriter Plus; the LaserWriter II series gives more uniform blacks. As with all xerographic printers, the edges of each character are a little rough, and printing can leave bits of stray toner on the paper. But a laser printer's tremendous flexibility with graphics and typefaces offsets these problems.

A LaserWriter's processing and paper-output speeds vary considerably depending on what's being printed. For ordinary text, PostScript itself is fairly fast, usually able to generate about five pages a minute, as long as the needed fonts are in the character cache. If the application program needs to add or change the fonts in the cache, the printer runs at about two or three pages a minute. Pages with complex graphics take still more time; in some cases a page may take an hour or more to print.

After the page image has been transferred to the page-image buffer, printing takes 12 seconds a page. For multiple copies of the same page, a LaserWriter runs at maximum speed after the first copy.

Paper designed for photocopying works better in laser printers than paper with a high cotton rag content. For offset masters, use 70-pound coated stock, such as Mead Offset Enamel or Shorewood Gloss, and start with the printing density control turned all the way up. Do not run Tyvek through a laser printer or copier; this synthetic material, which is widely used for tough, lightweight envelopes, will melt at the toner-fusing stage and jam the mechanism. (Note that neither Canon nor Apple recommends printing on both sides of the paper with a LaserWriter.)

Toner Cartridges

Both the LBP-CX and LBP-SX printing mechanisms rely on a disposable cartridge that contains critical components, such as the photoconductor drum and toner. When the toner runs out, or if the photoconductor becomes seriously scratched, you simply replace the cartridge. So far, Canon is the only manufacturer of toner cartridges for the LBP-CX or LBP-SX.

Because they contain so many intricate parts, cartridges are expensive. Under normal use, though, the drum is still good after the toner is exhausted, so a cartridge can simply be reloaded with new toner for about half the cost of a new one. Usually you exchange your old cartridge for a reloaded one. But is the reloaded cartridge on its first recharge or its tenth? For the best quality control, deal with a company that will reload and return your own cartridges.

Particulars: The LaserWriter and LaserWriter Plus

The LBP-CX printing mechanism in the LaserWriter and LaserWriter Plus is similar to the one in Canon's PC series of small photocopiers; it has been used by all the first truly low-cost laser printers. (LBP stands for Laser Beam Printer.) Designed for light-duty printing, the LBP-CX was rated for 3000 pages a month and a total lifetime of 100,000 pages, for a final cost of nearly ten cents per page. In practice, the mechanism has proven more robust; some units have already far exceeded these page limits.

The EP cartridges for the LBP-CX come in black, brown, and blue and contain enough toner for 3000 pages of average coverage; pages with large printed areas will reduce the page count. They are not, however, mechanically or functionally interchangeable with PC copier cartridges. In the copiers, black areas of the original are black in the copy (positive imaging or write white), whereas in the printers, the laser light writes the black areas of the finished page (negative imaging or write black). The LBP-CX incorporates a safety interlock for the laser beam, which is activated when the EP cartridge is removed, so the laser is active only when a photoconductor drum is present.

The paper cassette trays in LBP-CX-based printers, which look the same as but are not interchangeable with PC copier trays, can handle 16- to 21-pound (64 to 80 g/m^2) paper; the manual sheet feeder, 16- to 34-pound (60 to 128 g/m^2). The input tray holds 100 sheets, but the output tray holds only 20. Especially when the printer is running in a local area network, therefore, someone must check the paper trays regularly.

Paper is carried through the LBP-CX printing mechanism by a small plastic belt, which prevents printing within 0.35 inch along the right edge. Although the LBP-CX can otherwise print to within 0.2 inch of the left edge, 0.04 inch of the top, and down to the bottom edge, Apple's printing resource keeps a 0.42-inch border to square up margins and reduce memory consumed by the page buffer. Special software, however, will let you print to the LBP-CX's limits.

When you print multiple-page documents from most application programs, the stack comes out with the first page face up on the bottom. To avoid having to sort the pages once your document is printed, you can remove the output tray and let the paper drop 10 inches. The pages will flip over as they come out, landing first page first. (Some programs, such as MacWrite and Microsoft Word, can print from back to front.)

Several companies make paper-handling accessories for some Canon LBP-CX units—typically dual paper-feed trays (for letterhead and continuation sheets, for example) plus an envelope feeder. Most of these paper feeders are expensive.

Particulars: The LaserWriter IINT and IINTX

For the fastest printing of complex pages, the LaserWriter IINTX is the model to choose; it has a fast 68020 processor and can accept a lot of RAM and a SCSI port for an external hard disk drive to store font information. Both models also have an ADB (Apple Desk-Top Bus) port for additional flexibility. The printers can be connected through the ADB connector to receive information that is not part of PostScript instructions—information for controlling paper bins, for example.

Canon's LBP-SX printing mechanism in the LaserWriter IINT and IINTX has a longer lifetime—300,000 pages—than the LBP-CX as well as improved print quality, particularly smoother and denser blacks. The toner cartridges in the LBP-SX and LBP-CX are not the same, although they work in the same way.

The paper cassette tray in the LBP-SX can accept up to 200 sheets of 20-pound paper; larger paper bins are available and are simpler to install than in an LBP-CX printer. The manual paper feed will accept 8- to 36-pound paper. Optional cassettes are available for several paper sizes; an envelope cassette can hold 15 envelopes.

After printing, paper can come out either on top of the printer, face down, stacked in the order printed, or out the back of the printer, face up, stacked in reverse of the order printed.

A LaserWriter IINT can be upgraded to a IINTX by an Apple dealer.

OTHER LASER PRINTERS

Laser printers can be divided into four general categories:

- Models with a PostScript processor licensed by Adobe Systems
- Models with a clone of Adobe's PostScript processor
- Models with QuickDraw
- Models without PostScript or QuickDraw

PostScript Printers

Most PostScript printers licensed from Adobe include an AppleTalk port and require little more than plugging in to run with a Mac. Many models are the functional equivalent of a LaserWriter; the PS 810 from QMS (Mobile, AL), for example, uses the LBP-SX engine and performs somewhere in between the LaserWriter IINT and IINTX models. PostScript printers differ in several respects:

- Speed. Adobe has produced several versions of its software. The more recent ones are generally faster than earlier versions. Printers containing Adobe's PostScript processor with a 68020 CPU are also faster—about 20 percent faster than printers, such as the LaserWriter, with a 68000. For some PostScript printers, such as the Personal Pageprinter announced by IBM, the processor is housed on an accessory card inside the microcomputer instead of in the printer. Such processors send information to the printer more quickly than those connected through AppleTalk, but sharing the printer is more complicated.
- Configuration. So far, all of Adobe's announced PostScript processors—whether inside the microcomputer or inside the printer—incorporate a 68000 or 68020 CPU dedicated to PostScript processing. PostScript processing could be done by the computer's own CPU instead, but this would slow operation because the CPU would have to go through two computing steps: first calculating the PostScript instructions and then interpreting those instructions for the printer. This scheme would cost less, however, than a dedicated PostScript processor.
- Resolution and dot size. New PostScript printers with 400 and 600 dpi have appeared from Agfa-Gevaert (Mortsel, Belgium) and Varityper (East Hanover, NJ). The additional resolution makes for smoother images, but the dot size on these printers is not strikingly smaller than a LaserWriter's, so these printers are hardly substitutes for a typesetting machine (for which dot size is less than 1 mil).
- Image polarity. Some laser printers are positive printers, in which the photoconductor drum starts out black, and the laser writes the areas that will be white on the printed page; others, like the LaserWriter, are negative, in which the laser beam writes the areas that will be black. (Some printer

manufacturers define positive and negative with respect to the photoconductor drum, reversing the usage here.) Print quality differs subtly between these two printer types. With negative imaging, a character is built up from black dots; with positive imaging, characters are chiseled out of the black drum with white dots. Because of this difference, a one-pixel-wide line is practical on a negative printer. On positive printers, however, single-pixel lines tend to be obliterated, so these printers usually print lines at least two pixels wide; they do not print tiny characters well. PostScript processing details differ for each printer type.

PostScript Clones

Now that PostScript is in the public domain, many companies have announced their intention to develop a clone of Adobe's PostScript processor that will be able to accept PostScript commands. Because the clones won't be licensed from Adobe, they will probably cost less than Adobe's processor. Most clone makers claim that their products will run much faster as well, a claim that was untested when this was written.

There is no reason that a PostScript clone should not perform adequately, but because Adobe does not disclose details about how it encodes its proprietary fonts, clones might have trouble with Adobe fonts. In terms of font choice, this might not matter if enough quality fonts become available from other sources, but if what you are printing is to go to an Adobe-licensed PostScript typesetter, neither text nor graphics is likely to match up exactly. In any case, if you buy a printer with a PostScript clone, be aware that even if it passes casual tests for compatibility, it might fail tougher tests if you acquire new software that uses more PostScript functions.

QuickDraw Printers

QuickDraw is a set of routines based in the Macintosh ROM that creates the screen image as well as printed output on an ImageWriter. Two laser printers designed specifically to work with the Macintosh, Apple's LaserWriter IISC and General Computer's PLP (Personal Laser Printer), process QuickDraw instead of PostScript.

Although the two printers are based on the same idea, they operate differently, particularly in how they handle fonts. The LaserWriter IISC prints fonts as an ImageWriter does: It creates a font on paper from a screen font four times as large (a 12-point font on paper is scaled down from a 48-point screen font). General Computer's PLP contains outline fonts developed by Bitstream (Cambridge, MA), which, like PostScript outlines, can be scaled to different sizes. Neither printer can handle rotated type or condense a font to make it narrower, which PostScript printers can do easily.

The IISC and the PLP both connect to a Mac through the SCSI port. Both are designed to be single-user machines, rather than shared on a network, although General Computer has announced a networking accessory for the PLP. The LaserWriter IISC can be upgraded to a LaserWriter IINT or IINTX with PostScript, and General Computer has announced a PostScript upgrade for the PLP as well.

Because the QuickDraw laser printers are different from most Macintosh printers, application software—particularly graphics programs—may have trouble driving them; incompatibilities are likely to persist for some time. As these are worked out, the LaserWriter IISC will likely receive greater attention and support from software developers than the PLP, simply because it comes from Apple.

None of the Above

Other laser printers on the market are aimed at IBM PCs and interpret neither PostScript nor QuickDraw. The simplest, least expensive units cannot deal effectively with quality fonts or Macintosh graphics, although they can be useful for printing plain text files quickly. Hewlett-Packard's LaserJet is one of these; Phoenix Technologies sells printing resources for it and compatible printers. Some low-cost laser printers—including the Hewlett-Packard models—can be upgraded to handle PostScript, but the cost is usually the same or higher than it would be to buy a PostScript printer to begin with.

Coming Soon

Laser printers will reach a resolution of about 800 dpi, which appears to be the limit. Although laboratory experiments have created electrostatic images with more than 1000 dpi, the toning and fusing steps squash the toner particles, reducing the effective resolution. At 800 dpi, one page contains 59,840,000 bits, or about 7.5 MB, which places great demands on the control electronics and communication link between computer and printer.

OTHER PRINTING TECHNOLOGIES

Several companies have begun producing other kinds of printers, which, like laser printers, print an entire page at a time. Most of these also incorporate a xerographic drum with toner development but fewer moving parts. Arrays of light-emitting diodes (LEDs) or LCDs acting as shutters to a light source can create images on a photoconductor drum. In another approach, Delphax (Westwood, MA) offers a printer with a unique array of ion guns that squirt ions onto a photoconductor drum.

Magnetic printers record an image on a magnetic belt by means of an array of tiny recording heads. Toner particles, themselves containing magnets, adhere to the pattern recorded on the belt by the heads and are then transferred and fused to the paper. The magnetic image remains on the belt until it is erased, so many copies can be made from one image.

Color Printers

Developing medium-resolution color printers for microcomputers has proven difficult. Two technologies do work, although the results are often problematic. A color laser printer works like a monochrome one except that it has at least three separate toning steps, one each for yellow, magenta, and cyan; for the best results, a fourth step, black, may be added. Each step may have a separate photoconductor drum, or all the toning may take place on a single large-diameter drum.

Manufacturing and maintaining a color laser printer has been a major challenge; it is extremely difficult to get accurate registration and smooth blending of colors. Nevertheless, some prototypes have produced impressive results when they have been carefully adjusted by experienced technicians.

Medium-resolution color can also be produced by thermal transfer printers. These printers press a plastic sheet coated with color resin against the paper, and tiny heated pins melt the resin onto the paper; the resin-coated sheets are equivalent to a printer ribbon. The process is expensive because three or four 8½-by-11-inch sheets are used up for each full-color page. Dot size is much larger than it is with a laser printer, and the finished image has a shiny, glazed surface that may not withstand folding.

Ink-jet printers squirt tiny drops of ink onto the paper. Many current ink-jet printers are both low cost and low quality, but a new type with an array of ink jets will soon appear. The best will achieve medium resolution, and some will print in color. Such printers are quieter and capable of better results than impact dot-matrix color printers, but they will not compete directly with laser printers because of their relatively large ink-drop size.

PLOTTERS

Instead of printing dots, plotters draw lines by moving a pen directly on paper. Ballpoint, felt-tip, and liquid-ink pens produce a more even, continuous line than most dot-matrix printers, and many plotters offer a choice of pen colors. With suitable inks, a pen can write on transparent sheets for overhead projectors. Plotters come in many sizes; the largest can plot on bedsheet-sized paper, which is far bigger than what any printer can handle.

But plotters generate characters slowly and can't shade or fill in regions as well as dot-matrix printers can. As dot-matrix resolution has improved, plotters have lost ground, more and more becoming tools for specialists. There remains, nevertheless, a class of images best handled with plotters: large line drawings, particularly those in color. Plotters work best on structured images that consist mainly of lines; they work very poorly on bit-mapped images such as MacPaint pictures.

Some software drivers for plotters will not print bit-mapped pictures because the pen must be moved up and down repeatedly for each pixel, and this can prematurely wear out the plotter mechanism. The hardware-software package MacEnhancer, from Phoenix, includes plotter drivers that let you print directly from application programs; you can even assign pen colors to specific fill patterns in Macintosh graphics applications. Other companies that produce plotter drivers include Microspot (Maidstone, Kent, UK) and Mesa Graphics (Los Alamos, NM).

Sectio

Working with Macintosh

Three

What can you do with a Macintosh? What kind of programs are available today? What happens when things go wrong? This section surveys the applications and programming languages you will find for the Macintosh, introduces you to desktop publishing and computer communications, and provides advice on the many activities that are part of working with a Macintosh from day to day.

9: Words

Almost everyone who buys or uses a computer eventually does some form of word processing on it; in fact, many people use their machines for nothing else. With a computer, writing and editing anything from a letter to a book becomes easier, faster, and far less messy.

Instead of scrawling out drafts by hand or banging out innumerable versions on a typewriter, you type your text into the computer only once, changing words and phrases on the screen as you type. The computer allows you to move paragraphs, delete sections, and add new sections—all without having to type pages over and over. When you like what you've written, you can print a precisely formatted copy immediately.

At least that's the way it should be. A major challenge for word-processing program developers has been to give you on paper exactly what you see on the screen—not an easy task. With earlier word-processing systems, you could rarely be sure exactly where new pages of text began, much less see italics or headlines on the screen.

The Macintosh has changed all that. At last, what you see really is what you get.

WORD PROCESSORS

All Macintosh word-processing programs include traditional features such as word-wrap, basic formatting, cutting, copying, and so forth. In addition, most Macintosh word processors include features that were not commonly available before the Mac appeared, such as the ability to change the size and type style of characters and insert graphics into text.

Choosing a word processor depends on the features of a particular program and how they match your own preferences. It can also depend on which program your colleagues use; you can exchange files most easily if you are all using the same program.

The following brief guide to word-processing software emphasizes differences among programs. Because many programs are currently being revised, no attempt is made here to compare them in detail; refer to Macintosh magazines for more current information.

MacWrite

The first word processor for the Mac, MacWrite, version 5.0 (Claris, Mountain View, CA), is showing its age; overall, we do not recommend it. The program can deal with only one document at a time, and its formatting capabilities, designed for the ImageWriter, do not take good advantage of a LaserWriter or other PostScript printers. It does have one important feature: automatic repagination, in which page breaks are shown on screen at all times and are adjusted automatically as you type. A built-in spelling checker has also been added.

Because MacWrite was originally included in the price of a Macintosh, nearly everyone who needed to distribute text on a Mac disk—whether it consisted of letters or information about a program—used MacWrite's file format, secure in the knowledge that every Mac user had the program. Today the program is no longer included with Macs, but its file format remains common. Many other word processors can read this format; a few (for example, Microsoft Word, version 3.0, and WriteNow, version 2.0) can write it.

WriteNow

A medium-power word processor handy for business memos and reports, WriteNow, version 2.0 (pre-release, T/Maker, Palo Alto, CA), is easy to learn, has automatic repagination, and lets you open several documents at a time (how many depends on how much memory you have available) or write and edit two columns on screen at the same time. The program's built-in spelling checker works well. Its formatting capabilities are only fair, however; you cannot choose type in any point size or adjust the spacing within a paragraph independently of the spacing between paragraphs. Overall, the program is not suitable for complex documents. Unlike version 1.0, version 2.0 is supposed to support the arrow keys. WriteNow reads and writes files created by MacWrite and Microsoft rich text format (RTF).

LaserAuthor

LaserAuthor, version 1.0 (Firebird Licensees, Ramsey, NJ), was developed in England as MacAuthor. Its comprehensive formatting is based on style sheets—specifications for how a section of text will look on paper. An unusual feature is its ability to overtype any character with any other character. The program offers automatic repagination and particularly good search-and-replace features, but it lacks a built-in spelling checker. It reads but does not write files created by MacWrite and Acta, an outlining desk accessory (Symmetry, Mesa, AZ).

No page-makeup program can read LaserAuthor's formatted files, so using it might limit your options. On the other hand, LaserAuthor itself has some page-makeup features built in; you can, for example, see and edit text in multiple columns.

Microsoft Word

Word, version 3.0 for the Macintosh (Microsoft Corporation, Redmond, WA), which differs considerably from version 1.05 (and its MS-DOS namesake), is a powerful, complex word processor. Comprehensive formatting options, controlled by style sheets, can take full advantage of a laser printer. Built-in features include a good spelling checker, a rather awkward outliner, and a math function that lets you do calculations with any numbers displayed on the screen. The keyboard support is good; nearly all mouse functions have a keyboard equivalent. Many users, however, will be bothered by Word's lack of automatic repagination and its inconsistent handling of simple matters such as selecting words from the keyboard.

Nevertheless, Word is now the mostly widely supported comprehensive word processor for the Mac. PageMaker 2.0, for example, both reads and writes Word 3.0 files. And you can convert Word files to popular MS-DOS word-processing programs more easily than you can convert files from other Mac word processors. Word 3.0 reads and writes files from MacWrite, Word 1.05, MS-DOS Word, and Microsoft rich text format.

Microsoft Write

Microsoft Write, a simplified and cheaper version of Microsoft Word, has all essential word-processing functions, including spelling checking, but leaves out more sophisticated features such as style sheets and outlining; like Word, it also lacks automatic repagination. In essence, it contains the features of Word's Short Menus. Microsoft Write files are compatible with Word files. The program is an obvious choice if you don't need Word's power now but may move up to it in the future.

MindWrite

If you use outlines frequently, MindWrite, version 1.1 (Access Technology, South Natick, MA), is a word processor to try; it has the best integration of outlining into a word processor. Another useful feature is that any text you cut or copy is placed in an accumulating clipboard; you can retrieve any part of this clipboard for later use. MindWrite lacks a built-in spelling checker, however, and its formatting options are roughly equivalent to MacWrite's.

Microsoft Works

Microsoft Works, version 1.1, contains a spreadsheet, database, communication program, and word processor in a single package. The word processor is a medium-powered program with automatic repagination. The program is a good choice for those people who need its combination of applications. Because it has a built-in database, for example, Works is particularly effective for combining form letters with address lists stored in the database (mail merge). It lacks a spelling checker, but a separate spelling and punctuation checking program, WorksPlus Spell (Lundeen & Associates, Oakland, CA), can be added to it; this program also adds a glossary function (which allows you to store pieces of text you use often) and hyphenation.

Others to Come

As of early 1988, two other word-processing programs were being developed but had not yet been released: FullWrite Professional, an ambitious, comprehensive word processor with many page-makeup features (Ashton-Tate, Torrance, CA), and WordPerfect Mac, which will complement WordPerfect for IBM PCs, Data General, and Apple II computers (WordPerfect, Orem, UT).

WRITING AIDS

All writers need help, and today, there are almost as many electronic tools as printed ones to aid in organizing and polishing one's writing.

Outlines

Any word processor will let you create an outline. But word processors can show only a screenful of text at a time, usually only enough to see details of one part of your outline. Ideally, when you look at an outline, you want to see both the main headings *and* the details. Outline processors let you do this: With specific commands, you can collapse the outline so that only major headings show and then selectively open up portions of it to see successively more subheadings.

Some word processors have outlining built in. MindWrite's is excellent, and Microsoft Word's is functional. Acta is a good desk-accessory outliner that can work with any word processor. Sidekick (Borland, Scotts Valley, CA), a collection of desk accessories, includes an outliner as well. MORE (Living Videotext, Mountain View, CA), an application program originally designed as an outliner, also has a miscellany of other functions, such as creating organization charts.

Spelling Checkers

Spelling programs compare your text against a stored word list and flag all mismatches. Literal spellers use a word list containing only complete words; singular and plural forms are stored separately. Root-word spellers combine root words with many standard prefixes and suffixes; plurals are constructed by adding *-s* or *-es* to the root words. Root-word spellers take up much less storage space than literal spellers but will not catch such wrong words as *sheeps* unless they have an exception list.

Spellers cannot check for homonym mistakes (*peace* of cake for *piece* of cake) or for grammar, so you should never trust a speller to find all errors; you must proofread the final text. Spellers ought to find doubled words *(the the),* but many can't.

Now that most word processors come equipped with built-in spelling checkers, most users will not need to search out a separate one. Spellers are usually designed to work with specific word-processing programs; the one you choose must not disturb or be disturbed by the formatting information in your files.

Most spellers let you add words—such as a specialized vocabulary or names you use often—to their spelling lists. A good checker helps you with corrections, displaying close matches to mismatched words and saving you the trouble of looking up possible spellings. In general, the larger the spelling list, the more useful it is—but only up to a point: Immense spelling lists include many archaic and obscure words, so many typographical errors can be passed as valid words. The word *fro,* for example, is a legitimate word (as in *to and fro*), but it should be flagged as an error because it is most often a typo for the word *for.*

Thesauruses

Although less common than spelling checkers, electronic thesauruses are nearly as useful. Word Finder (Microlytics, East Rochester, NY) works much better than a paper thesaurus because you can look up not only synonyms but also synonyms of the synonyms. Word Finder is a desk accessory that works with most word processors.

Punctuation and Style

SmartQuotes (Oak Square Press, Allston, MA) lets you take advantage of some valuable features built into most Macintosh fonts: true opening and closing quotation marks (" ") and dashes (—), conventionally indicated by a double hyphen. The program, which can be installed as a desk accessory or INIT program, monitors your typing. Whenever you type a space followed by a ", it inserts a “ instead; if you type a " before a space, it inserts a ”. Two hyphens are replaced with a dash. These characters are available from the keyboard if you press Option or Shift-Option with the {[and hyphen keys, but SmartQuotes saves you the trouble of having to memorize these key combinations and keep track of whether to use an opening or closing quotation mark. A similar quotation feature is built into some application programs.

Programs that check writing style and grammar are still primitive. They look for a fixed set of common problems—stock phrases indicating wordy prose, doubled words, or sexist phrases. If you write *in view of the fact that,* the program might suggest *because.* Such programs, although simple-minded in design, can help clean up poor writing. A perceptive user can even extract value from a simple word-frequency analysis. In typical business writing, for example, the word *of* is more commonly used than *and*; good writing reverses the order.

Much good writing bends the rules, however; grammar programs won't improve Shakespeare. Nevertheless, combining a grammar program's analysis with common sense will usually help. You'll soon learn to avoid many simple grammatical errors, and then you can concentrate on making complicated errors that no software can detect.

Several writing style programs are available. Doug Clapp's Word Tools, version 1.0 (Aegis Development, Santa Monica, CA), is marred by an inadequate manual. Tools for Writers (Kinko's Academic Courseware Exchange, Santa Barbara, CA) does a better job and includes exercises to aid writers who are serious about improving their writing.

OTHER TOOLS

After you have created and stored dozens of word-processing files on your hard disk, how can you remember which file or files contained a particular piece of information? Sonar (Virginia Systems Software Services, Midlothian, VA) indexes the contents of text files for retrieval later. You can search for a word or a phrase, or you can specify something as complex as "the-words-*microcomputer-word-processing*-and-*Macintosh*-within-20-words-of-each-other-but-not-if-they-are-near-*IBM.*" Sonar displays the paragraph containing each found item. Version 4.5 can search MacWrite, WriteNow, Microsoft Word, MORE, Trapeze (a combination spreadsheet and page-makeup program from Data Tailor, Fort Worth, TX), and text-only files.

Several text editors run as Mac desk accessories, so you can use them while you are working in another program. They are much more useful than Note Pad, the desk accessory that comes with the Macintosh, and all of them create text files that can be read by other word-processing programs. With more features than the others, MiniWRITER, version 1.37 (Maitreya Designs, Goleta, CA), is the first choice. MockWrite, version 4 (CE Software, Des Moines, IA, and also as part of Sidekick from Borland, Scotts Valley, CA), also works well.

If a desk accessory is the smallest form of a word processor, hypertext represents word processing at its most grandiose. Unlike most word-processing tools, hypertext is not a writing aid but a reading aid—a means of finding related information fast. Hypertext links documents together according to references or common passages. While reading one document, you can display all the passages within it that refer to an earlier document or that are quoted in a later document. By pointing to one of the passages, you can instantly retrieve the complete document from which the passage was taken. Hypertext provides a comprehensive structure for revisions and comments, as well as a means of following concepts across multiple documents.

A full-scale hypertext system requires considerable computing resources and disk capacity. Two simplified hypertext products are already available for the Mac: Guide (OWL International, Bellevue, WA) and Apple's HyperCard. (See Chapter 11.) Guide is more suitable for business documents because its links between documents are based on text. HyperCard (version 1.1) builds up links based not on text but on the region of the video screen that a word occupies; this is a cumbersome technique that makes the program difficult to use.

A few programs have gone beyond assisting you with writing to the point of effectively doing the writing for you; you merely need to customize the output for your specific needs. Business letter packages contain form letters, such as thanks for payment, requests for payment, and demands for payment. WillWriter (Nolo Press, Berkeley, CA) supplies a thorough instruction manual and the boilerplate necessary for creating a valid will.

CREATING EQUATIONS

Most word-processing programs cannot cope with the specialized needs of scientific or mathematical text. To build formulas and equations, you need a library of symbol fonts (such as the font built into LaserWriters and most other PostScript printers) or the ability to create them; you also need to be able to pick up and move individual characters as well as define and move a group of symbols.

Among the standard word processors, Microsoft Word is the only one that supports equations. It does so through what's called a markup language: You surround the text and symbols of your equation with formatting commands; you can see how the final equation will look on screen or on paper only after processing. Both formatting commands and the substance of your equations are part of your word-processing document, so you can always search for any text within the equations.

Equations can also be treated as object graphics (defined, as in MacDraw, by specific attributes rather than as a bit map) to be pasted into the text of any standard word processor. Two programs that do this are the application Expressionist (Allan Bonadio Associates, San Francisco, CA) and the desk accessory MacEqn (Software for Recognition Technologies, Rochester, NY). Expressionist has a full set of tools for generating equations easily and can run as a desk accessory or as an application. MacEqn is a much smaller program with fewer features than Expressionist, but it can also create equations effectively and can be used along with your word processor. Because both programs treat equations as pieces of graphics, you cannot do a word-processing search for the contents of an equation.

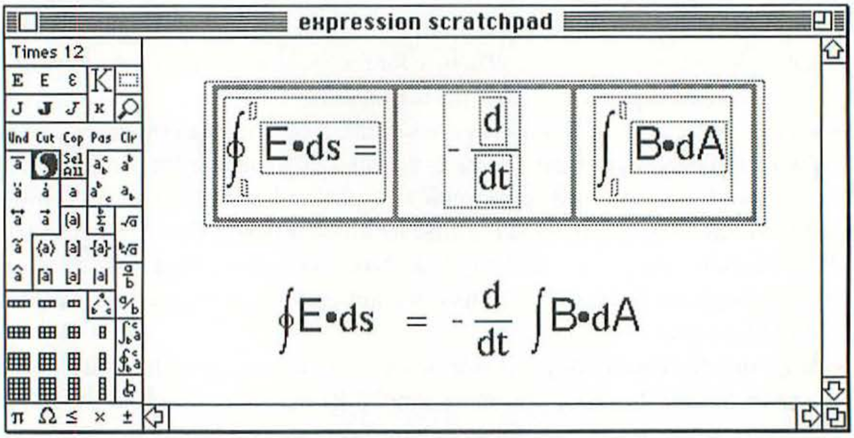

Expressionist designs graphics as object-oriented graphics.

The desk accessory ChemIntosh (SoftShell Company, Henrietta, NY) is a graphics editor specifically designed to create chemical structures. The structures are treated as object graphics and can be pasted into word-processing files.

FOREIGN-LANGUAGE WORD PROCESSORS

For European and other alphabetic languages written left to right and top to bottom, any Macintosh word processor will work. Standard Mac fonts include the diacritical marks many European languages need, although some accented characters are available only in lowercase. To generate accented capital letters, you need a modified font, or you can add them yourself with a font editor. (See Chapter 12.) Russian and Bulgarian require a Cyrillic font, which is available from several companies. Fonts for other scripts will appear in the future, or you can create your own.

For languages that read from right to left, standard word processors are unworkable. The program AlKaatib (Eastern Language Systems, Provo, UT) is for word processing right-to-left languages, including Arabic, Persian, and Urdu. It even handles the Arabic character set appropriately, changing a letter's shape according to its position at the beginning, middle, or end of a word, and forming ligatures (connected letters) as needed. HaKotev (also from Eastern Language Systems) and MouseWrite (Davka Corporation, Chicago, IL) are word processors for Hebrew.

Languages such as Chinese and Japanese pose other word-processing challenges. Chinese has an enormous number of characters—more than 74,000 in all—each representing a word; some 13,000 of these are needed for everyday speech and writing.

AlKaatib handles languages that are read from right to left.

Japanese is based on Chinese characters (kanji) but has in addition its own characters representing syllables (kana). Displaying Chinese characters clearly on a video monitor requires a lot of pixels. Roman characters are legible with an array of as few as 5 by 7 pixels; Chinese requires 24 by 24. For computers, Japan has a standardized character set, something like ASCII for languages based on the Roman alphabet; the standard for Chinese is likely to be the character set designated USMARC88.

FeiMa (Unisource Software, Cambridge, MA) and BrushWriter (Kaihin Technology, Singapore) are Chinese-language word processors. Both offer several ways to enter characters, including pinyin, today's standardized romanization of Chinese, and radicals (sets of strokes within a character).

EgWord is a Japanese kanji word processor written by Ergo Soft (Tokyo; available in North America from Qualitas Trading Co., Oakland, CA). You enter words from the keyboard in kana or romaji (roman alphabet), and the software offers the likely matching kanji characters, which you then select with the keyboard or the mouse.

Ecological Linguistics (Washington, DC) offers fonts for more than 30 other languages, mostly Asian, including Mongolian, Burmese, Telugu, and Thai.

TRANSLATION

Programs that assist translation from one language to another are beginning to appear. At this stage, these programs can only look up word equivalents and do rudimentary sentence analysis; all require a skilled human translator to produce a finished product. Nevertheless, these programs can speed up a translator's work, particularly when specialized vocabularies are involved. True translating programs that can give intelligible results on their own through syntactic and semantic analysis are still some years away.

10: Graphics

Pictures are the essence of the Macintosh. Being able to sketch or paint on the Mac's instantly erasable screen with the same finesse as on paper can do for graphic design what word processing can do for writing—make it easier, faster, and neater. Of course, good graphics, like good writing, require skill and patience as well as good tools; no fancy computer program will replace the hand and eye of a practiced artist, designer, or engineer. But the tools are getting better; today there are Macintosh graphics programs available for novices and professionals alike.

GRAPHICS SOFTWARE

Graphics programs come in many styles for varied uses. Which ones are right for you will depend on whether you do casual drawing and painting or publication-quality artwork and computer-aided design.

Traditional Mac graphics programs fall into two major categories: those that work with bit maps ("paint" programs) and those that use mathematical expressions to define lines and shapes (object-oriented "draw" programs). Paint programs deal with unstructured images, creating and storing pictures as collections of pixels. In a paint program, a line is a group of independent pixels, so you can change one pixel without affecting the others. These programs are good for quick free-form drawing; the tools available resemble pencils, paintbrushes, and spray paint.

In contrast, a draw program creates a line as a single unit. You can make the line shorter or longer or wider, but you are changing a line rather than individual pixels. Object-oriented programs are better for precise design work and drafting; the tools resemble a drafting kit with straight edges, French curves, and compasses.

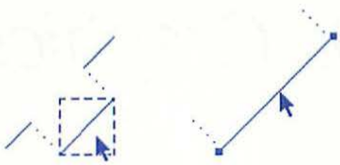

In a paint program (left), a line cannot be moved; only the pixels that are part of a line can be moved. In a draw program (right), a line, with small square "handles" to show that it has been selected, can be moved as a unit.

Object-oriented programs can be further divided into QuickDraw and PostScript varieties. QuickDraw programs generate graphics with the QuickDraw routines in the Macintosh ROM; QuickDraw images are then converted into PostScript for printing on a LaserWriter. PostScript programs create more sophisticated graphics by using PostScript directly without going through QuickDraw.

Paint programs usually work effectively with an ImageWriter or equivalent printer because the bit-mapped image on screen matches the printed output reasonably well. But paint images created at screen resolution do not take full advantage of the higher resolution of PostScript printers such as a LaserWriter. Everything else being equal, draw programs produce much better-looking results on a PostScript printer; a curve will look smooth rather than like a series of steps. Some bit-mapped graphics programs do support painting at 300 dots per inch (dpi) or more, although the sheer number of pixels at high resolutions means that you can spend an inordinate amount of time getting everything to look right. An object-oriented program is easier to use for 300-dpi images, provided it is suited to your subject matter.

Paint Programs

MacPaint set the standard for microcomputer graphics programs when it was first released with the Macintosh in 1984. It improved on earlier bit-mapped graphics programs for larger computers, offering striking performance despite the severe memory limitations of a 128 KB Macintosh. MacPaint, version 2.0 (Claris, Mountain View, CA), is the first significant improvement on the original program. It can open as many as nine pictures at a time and supports large screens. Each picture can be viewed at four magnifications. MacPaint, version 2.0, is better than FullPaint, version 1.0 (Ashton-Tate, Torrance, CA), which can open only four windows and only supports screens that are the size of a small Mac screen.

DeskPaint, version 1.0 (Zedcor, Tucson, AZ), which runs as a desk accessory, omits some traditional MacPaint features, such as spray painting, but adds others, such as erasers in more than one size. DeskPaint supports 300-dpi as well as 72-dpi images and can read and write TIFF (Tagged Image File Format) and MacPaint files. (See page 170.) DeskPaint requires less memory than other paint programs.

These paint programs can handle only black and white. To run them on a Mac II, you might need to set the monitor display to black and white. (Open the Control Panel desk accessory and choose Monitors; then choose "black and white" and "2 grays.")

To take advantage of the Mac II's color and gray palettes, paint programs that can handle color and grays are beginning to appear. Aimed at graphic artists, these programs have some distinctive features, including more elaborate tools than previous paint programs. Printing the resulting publication-quality output requires printers that can print in color and grays or typesetters with half-toning and color separation. (See Chapters 8 and 12.)

PixelPaint, version 1.0 (SuperMac, Mountain View, CA), works much like MacPaint but adds many enhancements to familiar tools. The paintbrush, for example, comes in different shapes and sizes and with ten special effects, including shadowing, whereby one brush stroke produces a primary stroke and a shadow as well. The spray paint can has speckling and eight other effects. With many tools, you can choose grading to make smooth changes from one color to another or from one gray to another.

ImageStudio, version 1.0 (Letraset, Paramus, NJ), deals with grays but not colors. It lets you draw freehand, but it lacks some common tools such as those for drawing rectangles and ellipses. Its main strengths lie in features for modifying and retouching images, particularly scanned photographs. You can change density and contrast, soften or sharpen borders, and solarize the image. You can even pick up and move objects in the scanned photo; if this is done with great care, the object will look as if it had been originally photographed where you put it. And there is a "finger" for smudging images, as you might smudge a real watercolor. ImageStudio pictures are best printed on a PostScript typesetter; the program's power does not come through on the coarse gray scale of a laser printer.

Although the capabilities of PixelPaint and ImageStudio overlap somewhat, the programs are best suited to different tasks: PixelPaint is better for creating images that are not meant to be realistic; ImageStudio is better for dealing with realistic images and images that were originally created in another medium and scanned into a Mac.

Object-oriented Programs

Draw, or object-oriented, programs have not received nearly as much interest from most Mac users as paint programs have because they do not lend themselves to free-form artwork; there is no pencil or brush that you can move in any direction. You cannot easily create a face or other images unless they can be constructed out of simple component parts—lines, squares, circles, and so on. You can, however, easily build and modify diagrams, such as a floor plan for a house, which is difficult to do with a paint program.

MacDraw was the first object-oriented graphics program for the Macintosh; it was originally published by Apple and is now sold by Claris. Few improvements have been made; although it is effective, version 1.9.5 is simple in concept and execution. For example, an object can be rotated only in ninety-degree increments; a drawing can be displayed only in three sizes. A new, updated version, MacDraw II, will appear in 1988.

MacDraft, version 1.2b (Innovative Data Designs, Concord, CA), was designed like MacDraw but includes many more features. You can see a drawing at eight different magnifications, rotate objects in one-degree increments, and attach numeric values to an object's dimensions. It is much better than MacDraw for careful drafting.

MacDraw and MacDraft share some design constraints because they are both based on QuickDraw rather than PostScript and are designed for output on an ImageWriter rather than a LaserWriter. Both programs will print on a LaserWriter, but object sizes and positioning are limited to fixed increments that correspond to screen pixels, and fill patterns available for shading or distinguishing objects are necessarily coarse.

In contrast, CricketDraw, version 1.1 (Cricket Software, Malvern, PA), is based directly on PostScript. As a result, it boasts many features that a QuickDraw-based graphics program cannot. CricketDraw can run text along a curve and decorate the text with drop shadows, for example, and it can fill objects with a graded tone from black to white. But although its printing repertoire is much more varied than MacDraw's or MacDraft's, CricketDraw is not suitable for executing something like a floor plan. Scale drawings are hard to create, and in any event, drawings in version 1.1 are limited to one page in size. Both MacDraw and MacDraft can create multipage drawings with flexible dimension scales.

None of these programs is suitable for producing polished commercial art, however. Adobe Illustrator, version 1.1 (Adobe Systems, Palo Alto, CA), can, in the hands of a skilled artist, produce finished artwork. It is not for doodlers or anyone without time to invest in mastering the program.

With Adobe Illustrator, a graphic artist can take a scanned image, such as this panda,

trace over it with curves,

and print a finished picture with a typesetter.

Illustrator normally produces finished art from templates; these templates can be created either through a bit-mapped graphics program or with a scanner. Illustrator reads in the image and supplies a set of curve-generating tools somewhat like a set of French curves. What you do, in effect, is trace over the bit-mapped image with curves. When you are done, you can discard the bit-mapped image. The curves are defined by Bezier mathematics, which are formulas for describing free-form curves, not merely fixed-shape curves like a parabola. For printing on a PostScript printer or typesetter, the printer's PostScript processor uses the mathematical formulas to construct the object at the printer's full resolution. Illustrator also supports halftoning and color separations, and Adobe has announced a new version that will add freehand drawing tools.

FreeHand (pre-release, Aldus, Seattle, WA) combines Bezier curves with tools for creating component elements. It can run in color on a Mac II and supports color separation.

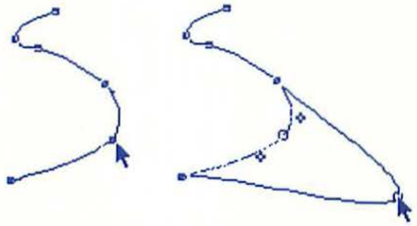

A Bezier curve can be fitted to any shape by stretching it like a rubber band (Aldus FreeHand).

Combining Paint and Draw

Since bit-mapped and object-oriented graphics programs each have different strengths, why not put them into a single program? Two programs do this. SuperPaint, version 2.0 (pre-release, Silicon Beach Software, San Diego, CA), and Canvas, version 1.0 (Deneba Software, Miami, FL), essentially combine the attributes of FullPaint with those of MacDraw. Both include Bezier curves, and SuperPaint adds other PostScript drawing capabilities.

Canvas can run as a desk accessory as well as an application program, but it is a little more difficult to learn than SuperPaint. Before you can paint with Canvas, you first have to create a frame and then paint inside it; you cannot paint anywhere on the screen. Both Canvas and SuperPaint have a valuable macro feature that lets you create new tools—to produce pentagons or starbursts, say—and use them as you would a MacPaint rectangle tool.

Both programs create bit maps at 300 dpi for a laser printer; Canvas can also create bit maps at up to 2540 dpi for typesetters. Because they combine both bit-mapped and object-oriented graphics, these programs offer a better introduction to Macintosh graphics than others.

COMPUTER-AIDED DESIGN

Computer-aided design (CAD) can be anything from planning a bookshelf or remodeling your basement to designing the space shuttle. MacDraw or MacDraft are adequate for your bookshelf, but they lack a CAD program's ability to create highly detailed engineering drawings suitable for specifying a crankshaft or an O-ring. In a CAD program, you do not normally build objects by selecting a tool, such as a line, and dragging out a line in the drawing area; instead, you select a line and then type in the exact length you want.

Since the output of a CAD program is usually instructions for producing a mechanical part, the programs are designed to parallel machine-shop practices. You can specify holes according to standard drill bit sizes instead of creating and sizing each one separately. You can select materials from a materials list; when you are done, the program can automatically compile a shopping list. It can also tell you how to get the sizes you need from standard stock sizes. If you needed plywood, for example, the program would figure out how many pieces you could get out of standard four-foot-by-eight-foot sheets, even allowing for the width of the cuts.

Two leading CAD programs for the Mac are MGM Station (Micro CAD/CAM, Los Angeles, CA) and VersaCAD (Huntington Beach, CA). Pegasys is a promising new program from IGC Technology (Walnut Creek, CA) that can read and write files generated by AutoCAD, the most popular CAD program on the IBM PC. Innovative Data Designs has announced Dreams, a family of CAD software.

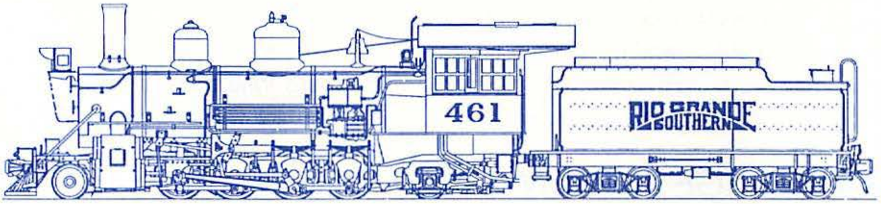

Output from Pegasys printed on a LaserWriter.

OTHER GRAPHICS PROGRAMS

For animation, VideoWorks II (MacroMind; published by Broderbund Software, San Rafael, CA), which works with images from virtually any paint program, is particularly effective. For example, you can create a foreground independent of the background and animate only the foreground. After you have strung together a sequence of images for animation, you can control the rate at which they are played back and add a sound track to accompany them.

Three-dimensional programs for the Macintosh were still embryonic in early 1988. Pro 3D (Enabling Technologies, Chicago, IL) and Super 3D (pre-release, Silicon Beach Software) show the essentials of how three-dimensional drawing programs work. As you draw, they display views of the image from three axes. You can change the point of view at will and make solid models with shading on the surfaces. But you pay a price for this power: Even on a Mac II, three-dimensional graphics programs run slowly (a graphics processor would help).

Once you start using graphics programs, you will build up a library of images and then have to keep track of them. Picture Base (Symmetry, Mesa, AZ) creates libraries of bit-mapped images that you can recall quickly. Curator (Solutions International, Montpelier, VT) does the same for common object-oriented formats.

MOVING GRAPHICS BETWEEN PROGRAMS

Most of the graphics programs described here store images in their own file format. Nevertheless, you can often transfer graphics between programs. Most bit-mapped graphics programs can read and write MacPaint files, and the newest programs can often read and write TIFF files. TIFF (Tagged Image File Format) is a bit-mapped graphics file format developed for scanners; it can handle grays and images of higher resolution than MacPaint files. Unfortunately, there are several variants of TIFF files, and they are not all compatible.

Object-oriented graphics programs can exchange information in several formats. Many programs can at least read MacDraw files, and the more flexible programs can read and write Apple's PICT and PIC2 QuickDraw files. In addition, object-oriented art or graphs—whether from draw, graphing, or spreadsheet programs—can usually be transferred to other object-oriented programs via the Clipboard or Scrapbook.

An object moved from an object-oriented program to a paint program loses its definition as an object and becomes a collection of pixels. Similarly, a bit-mapped picture that is moved to an object-oriented program can no longer be manipulated; the entire picture becomes a single object.

OTHER WAYS TO GET ART INTO YOUR MAC

You can get pictures into your Mac by means other than drawing or painting them yourself, including using clip art, tracing, overlays, and scanners. (Using scanners is probably the best way; see Chapter 12.)

Clip Art

In 1984, many companies rushed out MacPaint clip art, collections of pictures on disk something like the books of printable clip art. Much of this clip art is not very useful; how many times will you need a picture of Albert Einstein? And if you have access to a scanner, collections of bit-mapped images become superfluous. On the other hand, because of the detailed touch-ups necessary on some scanned images, some specialized MacPaint-based clip art can be useful. Decorative borders, for example, are hard to get exactly right with a scanner.

Today, clip art comes in object-oriented as well as bit-mapped formats, such as those of MacDraw and Illustrator. These collections are much more useful than earlier clip art because the images can be resized and changed without scaling distortion. A good example is MacAtlas (MicroMaps, Lambertville, NJ), a set of world and U.S. maps that you can cut and paste into other artwork.

MacAtlas is a collection of maps in object-oriented format that can be scaled without distortion.

Tracing

You can trace artwork with a digitizer tablet or with the mouse. For a clean image with either technique, you might not want to trace entire lines. Rather, you could trace the end points of each line or other critical feature and then use the tools of your graphics program to add the lines and curves.

A digitizer tablet has a penlike stylus. (See Chapter 5.) You can also use the mouse for tracing pictures, although it is not really designed for doing so. A mouse is difficult to position accurately and registers only relative position—the distance and direction you have moved, not where you started and stopped.

In a real pinch, you can use the mouse to trace from a transparent overlay on the screen. Photocopy what you want to trace onto a transparent sheet of acetate and paste it over the screen. To avoid parallax problems, don't move your head while tracing. If you do this more than a few times you will probably decide to get a scanner.

11: Business Programs

Microcomputers have become essential business tools thanks to software, not hardware. The Apple II succeeded not because of a unique hardware design—it was, in fact, inferior to many competing microcomputers—but because it ran VisiCalc, the first spreadsheet program. Today, most commercial software is aimed at businesses; after all, that's where the money is.

This chapter touches on some business software for the Macintosh. The information provided is hardly exhaustive; check magazines and user groups for more details.

SPREADSHEETS

Spreadsheet programs—powerful tools for calculation and projection—are the most popular business application for microcomputers. Organized like a ledger sheet, a spreadsheet program allows you to enter columns or rows of numbers, define relationships among them, and calculate results based on those relationships.

Suppose, for example, that you want to calculate the overall difference in cost between two new cars. One model has better fuel economy but is priced higher. Will this model be cheaper in the long run than one that consumes more fuel but has a lower price? To do this calculation without a computer, you would systematically enter each cost item on a sheet of ledger paper. To get a yearly total, you would have to make assumptions about the cost of gasoline, how many miles you expect to drive in a year, the interest rate on a car loan, and so on. Since you could not be sure how many miles a year you might drive, you would have to repeat the entire calculation several times for high, medium, and low estimates. The process would be laborious; even an accountant might stop calculating and start guessing after one or two repetitions.

A spreadsheet lets you do the same analysis more easily. You enter your values into the spreadsheet's columns and rows, and then, with a mathematical or logical formula, you link any one entry to any other entry. Once you enter the values, the computer does all the calculations. For a single calculation, a spreadsheet might not be any faster than a calculator. But for repeated calculations with changing values (variables), the spreadsheet shines. If you change any number—such as the number of miles you drive each year—the spreadsheet recalculates the total automatically. You can run dozens or hundreds of calculations painlessly. In so doing, you could learn the true cost of owning a car—where the money really goes and which costs are irrelevant to your comparison.

Worksheet1

	A	B	C	D	E
1	CAR COSTS	Honda Civic	Ferrari Testarossa		
2	Price	$6,704	$102,500		
3	Down payment	$1,676	$25,625		
4	Loan value	$5,028	$76,875		
5	Monthly payment	$167	$2,553		
6	1st Yr cost	$3,680	$56,265		
7	2nd-3rd Yr cost	$4,008	$61,280		
8	Fixed costs	$800	$6,000		
9	Miles per gallon	32	12		
10	Gas cost	$1.20	$1.20		
11	Miles per year	12,000	12,000		
12	Running costs	$450	$1,200		
13	3 yr total	$8,938	$124,746		
14	Salvage value	$4,200	$90,000		
15	3 Yr cost	$4,738	$34,746		
16					

Microsoft Excel shows you the true cost of owning a Ferrari.

Spreadsheets can give you the edge you need to plan ahead in your business, too. For example, you can construct a spreadsheet model of your income and expenses, using formulas to add tax and shipping charges automatically to each purchased item and calculate discounts for volume purchases.

Spreadsheets are not only for calculations, though. If you need to produce an organization chart or other text arranged in columns and rows, a spreadsheet often works much better than a word processor. Whereas moving an entry from one portion of a table to another is difficult with most word processors, it is easy with a spreadsheet. (Most spreadsheets distinguish between text and numbers by the characters you type, but you can have them treat a number as text; refer to the manual for your spreadsheet.) When you have finished, you can copy the table to the Clipboard and then paste it into a word-processing document.

All spreadsheet programs can perform ordinary spreadsheet functions, but they differ in many details. You should probably choose one spreadsheet program and stick with it. You can move information between some programs, but you will often lose many specific features. (If you need to move data to or from an IBM PC, see Chapter 21.)

Comparing Spreadsheets

The Macintosh spreadsheet market is dominated by Microsoft Excel. Nearly all the support material for Macintosh spreadsheet users—books, templates, and training courses—covers Microsoft Excel. The program is excellent, but other spreadsheet programs (described below in alphabetical order) perform at least adequately, and some of them have distinctive features.

MacCalc, version 1.2 (Bravo Technologies, Gilroy, CA), is a straightforward spreadsheet without any graphing features. Its main advantage over Microsoft Excel is its much lower price.

MacPlan is a desk-accessory spreadsheet, a component of Sidekick, version 2 (Borland International, Scotts Valley, CA). Despite the necessary restrictions of a desk accessory, the program has a surprisingly complete set of features, including simple graphing, but the spreadsheet is limited to 50 rows by 20 columns. If your needs are modest—keeping track of small expense accounts, for example—MacPlan may be all the spreadsheet you need. MacPlan has two severe faults: It lacks an Undo command, and it does not warn you to save your data before closing a spreadsheet window.

Microsoft Excel, version 1.06 (Microsoft Corporation, Redmond, WA), is among the most powerful spreadsheets developed for computers thus far; it includes comprehensive graphing capabilities and simple database functions. You can create macros by assigning a series of commands and spreadsheet entries to a simple keystroke combination. You can also define your own functions and consolidate multiple spreadsheets.

Microsoft Works, version 1.1, a combination package, contains a spreadsheet and is a good choice for those who can make use of its other parts (graphing, word processing, database, and communications).

Modern Jazz (pre-release, Lotus Development, Cambridge, MA), the new version of Jazz, is an integrated package with a spreadsheet and graphing, word processing, communications, and a database. Modern Jazz's outstanding feature is its live links to the word processor. When you paste a spreadsheet into a text document and then change the spreadsheet, the document is automatically changed as well.

Ragtime, version 1.1 (Orange Micro, Anaheim, CA), boasts page-makeup features in addition to its spreadsheet functions. It is therefore particularly suited to preparing financial reports or other publications that contain spreadsheet data. Information from other spreadsheet programs can be pasted into separate page-makeup programs, but with Ragtime, everything is contained in a single program.

Trapeze, version 2.0 (Access Technology, South Natick, MA), breaks away from traditional spreadsheet design. In Trapeze you can create spreadsheets in the standard row-and-column format as well as place cells on a page anywhere you want—and fill in the rest with graphs and text. Unfortunately, Trapeze's Undo command is only partially functional, and the program has only limited ability to export and import information to and from other programs.

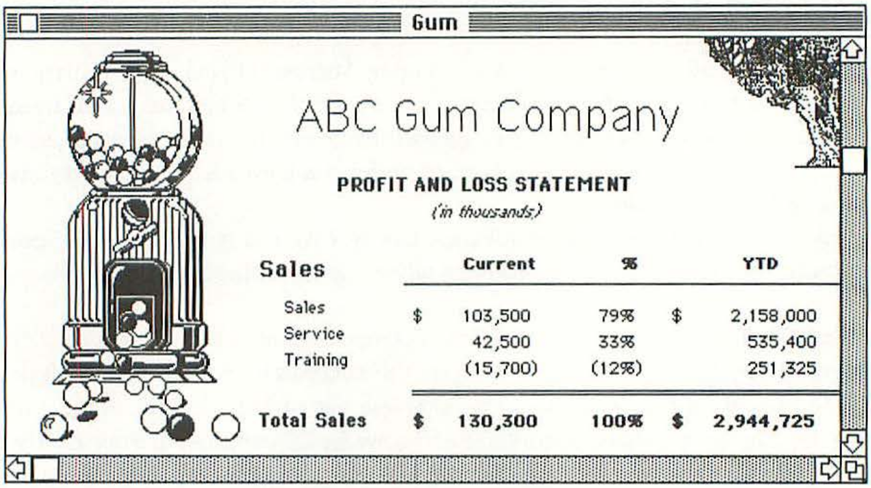

In Trapeze a spreadsheet layout can be more flexible than a grid of rows and columns.

CHARTS AND GRAPHS

A spreadsheet generates tables of numbers, such as those that appear in annual reports and statistical summaries. But a table may not be the most effective way of communicating the meaning behind the numbers.

People can usually recognize and interpret graphics much more easily than they can rows and columns of figures. For the same reason that the icons in the Macintosh Finder work more effectively than the massed text that passes for file directories on other computers, graphs and charts work much better for spotting trends or making comparisons.

You could use programs like MacPaint or MacDraw to create charts by hand, but to enter numbers and turn them automatically into charts, you need a graphing program. On the Macintosh, most graphing programs are part of a spreadsheet program; they plot numbers from the spreadsheet. Even if you don't need its computational power, the spreadsheet provides a useful vehicle for entering and checking your numbers before graphing them.

For more graphing features than spreadsheet programs provide, try Cricket Graph (Cricket Software, Malvern, PA). It can, for example, fit curves through the points plotted on a graph. Cricket Pict-O-Graph is a specialized graphing program that creates pictograms, graphs that display data with icons and pictures.

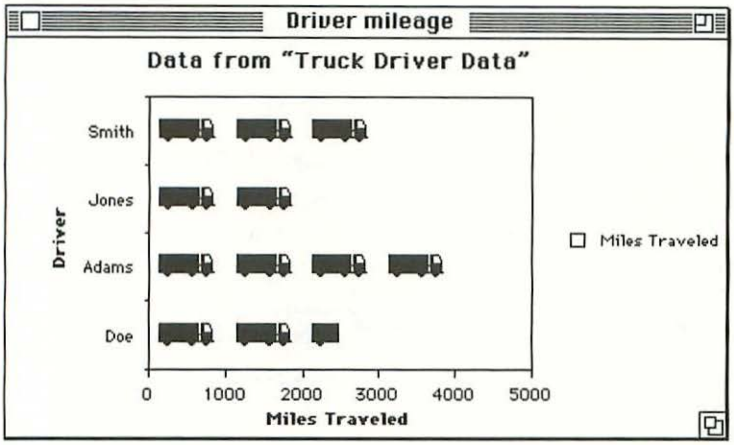

Cricket Pict-O-Graph draws graphs with pictures.

DataScan (ISS Design, Plainfield, NJ) is an unusual program that can transform scanned images of existing graphs from virtually any source—whether hand drawn or taken from a newspaper—into clean Macintosh images that can be placed in a page-makeup or other type of program. DataScan comes with many tools for cleaning up scanned graphs, including data interpolation and noise filtering for ignoring stray pixels in the image.

SOLVING EQUATIONS AND MODELING

Although all spreadsheet programs can perform common mathematical calculations, none are really designed to solve equations. Eureka: The Solver (Borland International, Scotts Valley, CA) can solve a wide range of equations, find the maximum and minimum values of functions, and plot the results. It can also generate reports, complete with graphs, as MacWrite files.

Spreadsheets can also help you track any situation that can be described in mathematical terms. They are designed to describe a particular instant, however; they have trouble modeling situations where the variables change over time.

STELLA (High Performance Systems, Lyme, NH) is a program designed for dynamic modeling. With it, you could simulate, for example, a production line in which manufacturing time and available resources limit output. STELLA can forecast the process through time by animating the model itself on screen or plotting it as a graph with time along one axis.

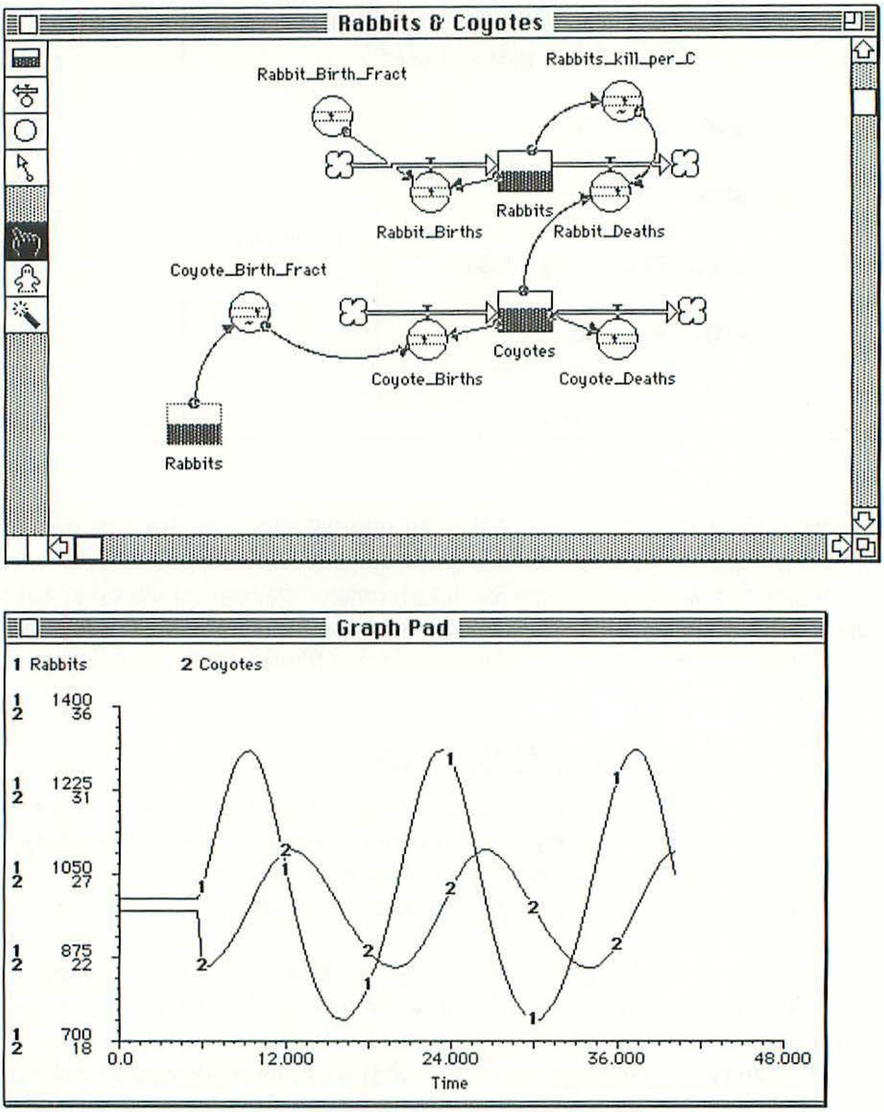

STELLA can display the struggle between predator and prey as an animated diagram (above) or as a graph.

DATABASES

A database is an organized collection of information, or data; for example, an address book is a database. A specific database file is made up of records (such as a complete address), each of which is made up of fields (name, street address, zip code, and so on). You index a database by field: Most address books are indexed by name, as are telephone directories.

One reason computerized databases are so powerful is that you can change the index field; if you are preparing a mailing, you can index by zip codes instead of names. You can also index according to more than one field. And you can select only those records with fields that meet specific criteria; for example, you can request records for all the Jacksons who live in Chicago.

A database program lets you enter, sort, update, find, and print your data quickly and easily. Although the idea of a database is simple, database programs are the most complex of common microcomputer applications. The most powerful go far beyond simple functions, letting you define relationships among data entries, perform statistical and accounting calculations, extract and reorganize the data, and much more.

Choosing a Database

Traditionally, the problem for database users has been deciding how much power to buy; the more powerful the database program, the more difficult it has been to use. The Macintosh interface has helped, but even though Mac database programs are easier to use than their equivalents on other computers, programs boasting great power still require considerable effort to master.

When shopping for a database program, consider these features:

- Multiple field types. The simplest databases use only text fields. But if a program recognizes only text and not numbers, it cannot do any calculations on fields. Nearly all programs let you define text and numeric fields, and many also include date fields for time-dependent financial calculations and logic fields for true-false information.
- Editing during data entry. You should be able to edit any text entry just as you do in a word processor—double-click to select a word, shift-click to extend a selection, and so on. You should be able to select a portion of any entry and cut, copy, or paste it to any other entry (pasting a text entry into a purely numeric field doesn't make sense, of course). For convenience, you should also be able to create a custom form on screen for entering data.
- Sorting. The program should let you select a field to sort by. When sorting a field alphabetically, the program should sort letters regardless of whether they are uppercase or lowercase (dictionary sort). For either alphabetic or numeric sorting, you should be able to choose whether you want to sort from lowest to highest or vice versa.
- Queries and reports. You should be able to ask for all records meeting certain criteria—for example, every salesperson who sold more than 500 widgets in a month.
- Flexible output forms. A database should let you organize and print out your data any way you want. You should be able to position and size each record with the mouse so that, for example, you could change quickly from an address directory to mailing labels. The reports should have a mode for printing on standard mailing labels.

- Print merge. It should be easy to combine information in the database with text from a word processor for form letters. More generally, you should be able to use the Clipboard or Scrapbook to move information between the database and other programs.

Other features are useful, but, depending on your needs, they may not be essential:

- Graphics. Some databases can store pictures and graphs as a field; the pictures might come from MacPaint or another program.
- Forms. A flexible database can create forms as well as mimic existing ones, such as tax forms. The Internal Revenue Service will accept the forms for tax returns printed by specific programs.
- Calculating on fields. Although they are unnecessary for address books, calculations are needed for many jobs, such as managing sales information and inventories.
- Variable-length fields. Most traditional databases accept only fields of fixed length, so you can't squeeze in a little extra information or annotate specific records. Most Macintosh database programs use variable-length fields, which let you add notes anywhere you want. There is a penalty, however: It usually takes more time to retrieve variable-length data than fixed-length data.
- Searching all fields. When you can't remember where you entered a specific item, the program should allow you to search all fields to find it.
- Indexing all words. If you routinely create records containing fields with more than one word, you might want to index all words in a field for retrieval later.
- Sloppy-search. Suppose you can't remember how to spell a certain entry. A sloppy-search function lets you type in something resembling what you want, and it will find close matches. You can often do this with a fragmentary search, using the letter(s) you can remember, for example; however, many database programs search only at the beginning of a field. A sloppy-search system should automatically search for homophonous sounds: If you enter *Brown,* the program should also look for *Braun.*
- Information interchange. Some database information is structured much like a spreadsheet. If you use both types of programs, you might want a database that works smoothly with a spreadsheet and lets you move information back and forth with a minimum of fuss. Integrated packages can tightly bind together the spreadsheet and database.
- Power. Will the software grow with your needs? In general, you should buy a program with a little more power than you think you need, rather than one you might outgrow soon. Some programs will create compatible data files that you can use with more advanced programs. To do a simple task quickly, consider a filecard database.

- Networks. Several companies offer sophisticated database programs for network operation; such programs permit several users to work simultaneously with a file. (See Chapter 20.)

Database Organization

Database programs follow three major models for organizing information. Flat-file databases store all data in a simple sequence for each record; the data can be represented in a grid, with the records as rows and the fields as columns. Hierarchical databases arrange fields or, sometimes, entire databases, in a tree-like structure; access to any particular field follows a branching path. Relational databases use fields common to several files to maintain relationships you have defined among them. The most powerful of these programs can maintain multiple relationships simultaneously.

Database Programs

At the high end of the market are Omnis 3 Plus, Double Helix II, and 4th Dimension. Omnis 3 Plus is a hierarchical database; the other two are relational databases. All three support networks. All three are complex products that take considerable time to master. In many businesses, a data-management specialist might develop the company's own applications from one of these programs, thereby sparing everybody else the hassle.

Omnis 3, Double Helix II, and 4th Dimension have developed in completely different directions. For most users, choosing among them is more a question of design philosophy than of specific features. Omnis 3 Plus, version 3.24 (Blyth Software, San Mateo, CA), was originally developed before the Macintosh appeared. Although many Macintosh features have now been built into the program, it still shows its pre-Mac origins; it is the least Mac-like of the three high-end programs. It cannot, for example, store graphics in a field or use fonts and type styles to differentiate data on screen, although it can print reports with different fonts. Blyth Software has developed a new database program, called QUARTZ, that adopts a more graphical interface, but QUARTZ so far runs only on IBM PCs under Microsoft Windows.

Double Helix II (Odesta Corporation, Northbrook, IL) is at the other extreme: If Omnis 3 Plus is not sufficiently Mac-like, Double Helix II is almost too Mac-like. With Double Helix II, you manipulate icons for everything, including structuring your records and doing calculations. To set up a database procedure, you organize icons (called tiles) into a flowchart; when you are done, the flowchart is the procedure. Double Helix II is so different from other databases that traditionalists will find it hard to learn, but for some Mac enthusiasts, it is the most natural database yet. Odesta also publishes Helix VMS, a program for Digital Equipment Corporation (DEC) VAX minicomputers that can turn a VAX into a database source for Macs connected to it.

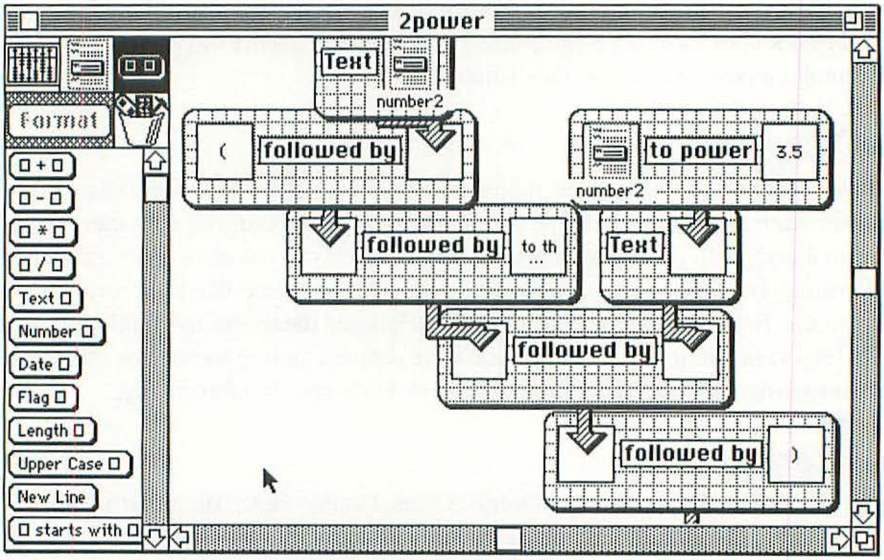

Programming $2^{3.5}$ in Double Helix. Calculation "tiles" are shown on the left.

Acius' 4th Dimension, version 1.0 (Cupertino, CA), works in a Mac-like way without going whole hog with icons. It has more features than any other Mac database program, and it supports serial port communications, graphing, and a complete programming language, so you can create your own procedures. But in its initial release, 4th Dimension runs very slowly and cannot be recommended for anything except Mac IIs.

Somewhat less comprehensive than the three high-end programs, Ashton-Tate's dBASE Mac (Torrance, CA) is a relational database with a simple procedural language. It does not support multiple users on a network. Despite its name, the program is a fully Mac-like database that is unrelated to Ashton-Tate's popular dBASE III PLUS for the IBM PC. The Mac product can read files generated by dBASE III PLUS, but it cannot run programs written in the procedural language of dBASE III PLUS.

Other products for the Mac do run dBASE III PLUS programs and use its file format as well. Such programs can take advantage of the many auxiliary products developed for dBASE III PLUS, including complete accounting packages written in its procedural language. FoxBase/Mac (pre-release, Fox Software, Perrysburg, OH) is a Macintosh database program that not only reads and writes dBASE III PLUS files but can also run dBASE programs. FoxBase has a Macintosh-style interface as well as a dBASE-style command-line interface. Nantucket (Los Angeles, CA) publishes McMax, which is another dBASE III PLUS-like product for the Mac. McMax does not work like a Mac program; instead, it faithfully replicates dBASE III's obscure interface as if it were running on an IBM PC. These dBASE-style products are relatively fast, partly because they have fixed-length fields and records.

Many people do not need a high-powered database with elaborate features. Reflex Plus (Borland International, Scotts Valley, CA) for the Mac is an attractive medium-power relational database. To keep simple lists, without relationships between the records, FileMaker Plus (Nashoba Systems, Foster City, CA) is a smooth, easy-to-use program that indexes every word in every field.

Finally, for very specific needs, there are three specialized database programs from Satori Software (Seattle, WA). Project Billing and Legal Billing manage time and resource billings for consultants and attorneys, respectively. Bulk Mailer organizes large address lists for mass mailings.

HYPERTEXT AND HYPERCARD

Although hardly a traditional database, Hypertext is nonetheless a data-handling system. Based on an idea by computer visionary Ted Nelson, hypertext gives you the ability to link and follow related information from one document (or card) to another—fast.

In classic hypertext, the links are based on text; if you move text, the link moves with it. Guide (OWL International, Bellevue, WA) follows this model; it cross-references documents of any length. You can scroll through each document as in a word-processing program. When you come upon, say, a name in one document that is also referred to in another, the mouse pointer changes shape, and the other document appears on screen at precisely the place where the name is mentioned. Thus, you could very quickly pull up all the letters mentioning a Ms. Alexander, no matter how long or how diverse the original correspondence might be.

In HyperCard, version 1.1, from Apple, the data consist of text and images on cards. The relationships linking the cards are created along with them; you can move from one card to another following related ideas. For example, an opening card might be a building floor plan: As you click on specific offices, new cards display the names and pictures of the occupants. Then, if you click on a name, you might see that person's employment record.

Unlike Guide, the links in HyperCard are principally graphical: You click on regions of a card to follow a link. Although a region can contain text, it is defined graphically and remains in place even if the text is moved out of it. HyperCard's displays are always cards, with one card showing at a time. Each card is the size of a small Mac screen, no matter how large the screen you are using. Thus, HyperCard is most suitable for items that fit naturally on a three-by-five card; the program cannot handle large documents well. It does, however, support sound.

Programs and data for HyperCard, called stacks or stackware (after a stack of cards), are available from many sources. Stackware uses up considerable storage space; a hard disk drive is virtually essential. Storing fifteen seconds of sound, for example, requires more than 500 KB. Early stackware products have displayed some imagination, but few are worth buying yet. More interesting stackware may arrive as programmers begin adding features, such as the ability to display more than one card at a time or to scroll records on the screen.

ACCOUNTING PACKAGES

Accounting programs are based on a database that keeps track of people owing you money, people you owe money to, and how much. Most programs also have special provisions that let you handle taxes. A simple accounting program—suitable for households or tiny businesses—is usually self-contained: It tracks money in (accounts receivable) and out (accounts payable) and keeps an overall summary (general ledger). More complex packages, suitable for small to large businesses, often come in several parts. Accounts receivable, accounts payable, payroll, inventory, and other components are sold separately, along with the general ledger package that holds them all together.

Both general-purpose and highly specialized accounting programs are available for the Mac; many can be modified to meet your needs. Particularly if you are seeking an accounting package for your business, it is a good idea to have a qualified accountant help you choose. See if the program will work with the other software you use. Can you generate graphs from the data? Or move the data to a spreadsheet to make financial projections? Can the program grow with your company? Can your accountants use your files on their computers to do audits?

To keep track of home accounts or a tiny business, MacMoney (Survivor Software, Inglewood, CA) is one alternative. It can track up to 6000 transactions in a year and print checks for you.

One specialized accounting program that meets a common need is MacInTax (SoftView, Camarillo, CA). The 1987 edition can display 30 federal tax forms and schedules on screen. After you enter raw data, such as your gross salary and taxes withheld, the program computes your taxes, fills in the remaining entries, and prints out the results as a tax form on plain paper. The final forms, whether printed on an ImageWriter or a LaserWriter, are accepted by the Internal Revenue Service. SoftView also has a program for computing California state taxes.

PRESENTATION GRAPHICS

The staple of business meetings, presentation graphics are designed to make a few points clearly with large type and uncluttered graphics. They can be either projected on a screen or handed out on paper. For projection, the graphics can be printed on a laser printer and photocopied onto acetate for overhead projectors or photographed on 35-mm film for slides. With suitable software, a slide maker can produce presentation graphics in color. (See Chapter 22.)

Many programs can generate presentation graphics: draw and paint programs, word processors, and page-makeup programs. You simply need to create output with no more than ten lines on a page; a simple boldface font is best. MORE, an outlining program from Living Videotext (Mountain View, CA), includes some presentation features, such as the ability to turn an outline into an organization chart. But for the greatest flexibility, particularly if you have to do a lot of presentations, try a program designed specifically for presentation graphics, such as PowerPoint (Microsoft Corporation, Redmond, WA) or the more elaborate Cricket Presents (pre-release, Cricket Software, Malvern, PA).

FLOWCHARTS

Flowcharts, often used for planning, can be created with any draw-type graphics program, but if you need more than a few, you should get a flowchart program. Design (Meta Software, Cambridge, MA) creates flowcharts easily; when you alter your chart, you can drag items, and any connected arrows follow along. Design supports all common flowchart symbols and conventions, and it lets you create your own shapes or paste in a picture created by another program.

PROJECT MANAGEMENT

MacProject II, from Claris, deals with time and resource allocations. You can examine time, personnel, budget, and material needs for a complex project and see which steps are the most critical to the project's success.

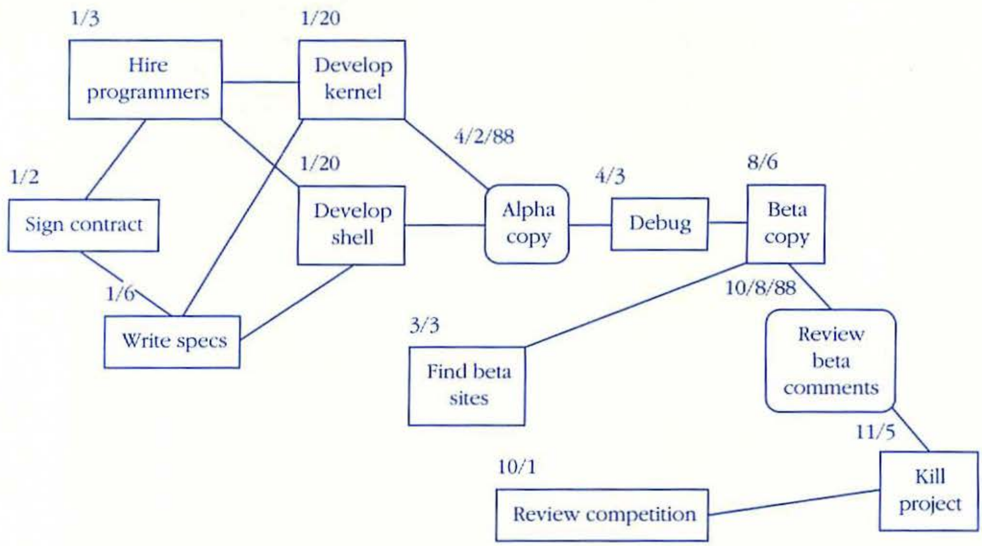

MacProject might help you avoid some project-management pitfalls.

MacProject and related programs are useful mainly for well-defined projects whose parameters can be specified precisely. For the chaos that dominates planning in most organizations, these programs are less effective; they don't handle fuzzy information and goals well. The more enthusiastic business school graduates argue that a project-management program should impose order on the chaos: Let the master plan, as expressed by the software, tell everyone what to do and when to do it. This crowd may be right or merely naive, depending on the situation—or your point of view.

TIME SCHEDULING

Nearly all of the many calendar programs for the Macintosh are impractical, for reasons explained in Chapter 16.

For major scheduling jobs, such as in a medical or legal office, Front Desk, version 1.0 (Layered, Boston, MA), offers comprehensive functions. It can track the schedules of up to 15 persons as it records billing time and canceled, missed, or rescheduled appointments. You can set up standard appointment lengths, block out holidays, and view a week's worth of appointments for five people at once on one screen. Front Desk even copes with common scheduling abuses, such as double-booking in a doctor's office.

"EXPERT" SYSTEMS

In the years to come, a valid class of "expert" system programs will appear—programs that can store knowledge based on specific, well-defined subjects. These programs will be able to analyze a particular situation and offer suggestions or predict potential problems.

Some already-available programs claim to offer expert advice. For example, several try to tell you how to negotiate or close a sale; you answer a series of questions about the person you are dealing with, and the program generates several pages of advice. But there's much less to these programs than meets the eye. You must know your "adversary" well to fill out the questionnaire accurately, and if you know the person that well, you probably don't need the program to begin with, even assuming a valid analysis—a dubious proposition.

12: Desktop Publishing

Publishing—transforming ideas, words, and pictures into a form suitable for distribution—is an art. An aesthetically pleasing publication that succeeds in getting its message across represents the skillful integration of text, design, and printing elements. In the past, each phase of publication was the domain of specialists with specialized tools, but microcomputers, particularly the Macintosh, have put many publishing tools into the hands of individuals and opened the way for a new art: desktop publishing.

Most of the steps in desktop publishing—choosing type and preparing text for typesetting, producing line drawings and photographic images, laying out and pasting up pages, and readying the result for final printing—are the same as those in traditional publishing; it is the technology that is different. Unfortunately, the technologies have changed faster than the elements they were invented to handle; too often, this results in awkward compromises on the printed page. Typefaces, for example, were originally developed for sculptured media; letters were chiseled into stone or molded from metal. Adapted for the pixels of a computer screen or printer, many typefaces have simply become clumsy renditions of once-elegant designs.

Graphic design, too, can suffer. Microcomputers have made processing graphics almost as accessible as simple typing. Will one million Macintoshes sold mean ten million bad pictures? Will new page-makeup programs mean reams of badly designed pages? Not everyone is an expert typesetter or designer, and new technologies at our fingertips won't change that.

TYPEFACES AND FONTS

A typeface is a named type design, such as Helvetica or Palatino. A font is a complete set of characters in a given typeface and size, measured in points. An American and British printer's point equals 0.013837 inch, or about 1/72 inch. A European, or Didot, point is somewhat larger—0.3759 mm, or 1/68 inch.

Three kinds of fonts exist for the Macintosh:

- Screen fonts, designed for the screen and for printing on an ImageWriter or LaserWriter IISC.
- PostScript fonts in outline form, designed for the PostScript LaserWriters and other PostScript printers and typesetting machines.
- Fonts for the screen that match PostScript outline fonts but are not designed to be printed. (See Chapter 8.)

The point sizes given in the style menus of many Macintosh programs are close but not identical to standard printer's point measurements. The Macintosh screen fonts and the ImageWriters use 72 points to the inch.

Screen Fonts

Screen fonts are designed to produce good results on the screen as well as on the ImageWriters and the LaserWriter IISC. These typefaces come in specific sizes, which are stored as bit maps of each character in the font information portion of the System file. To be available for the screen and for printing, each size of each typeface you intend to use must be installed either in the System file or in an application program. Fonts installed in the System file are available from all programs; the fonts installed within an application are available only from that application.

To install or remove a screen font from the System file, use the Font/DA Mover program supplied by Apple or an equivalent program. Your installation choices will include application programs.

If you choose to work with a size that is not installed, Macintosh system software will create it on the screen by scaling an installed size up or down. Some scaling factors work better than others. Scaling a 24-point font down to 12-point size, for example, generally works well because it is a simple matter of omitting every other dot horizontally and vertically, but scaling down to 17-point size leaves rough edges. (See Chapter 8 to learn how fonts are scaled for printing on an ImageWriter and the LaserWriter IISC.)

Installing every possible size is impractical because the fonts would take up too much storage space. You should, however, avoid scaling typefaces for printing. If you cannot find the font size you need among Apple's screen fonts, check with independent companies; many of them supply the sizes Apple does not, and so scaling is becoming almost unnecessary.

The proportions of the letters in some screen fonts have been fudged for legibility. Most often, the x-height (the height of the lowercase letter *x*) is taller than it should be in some small type sizes because the screen font's designers have few pixels to work with; for example, some 12-point sizes use the x-height of a 14-point font. And in some fonts, such as Geneva, line weights in the larger sizes have been kept thin so that they look correct when scaled down for Best mode printing on the ImageWriter or ImageWriter II.

You can choose from hundreds of screen fonts. Many are designed only for display printing (18-point and larger) because neither the screen nor the ImageWriters have enough pixels to resolve fine serifs or curlicues in smaller sizes. Each screen font can contain up to 255 characters. Here are some samples from CasadyWare (Carmel, CA):

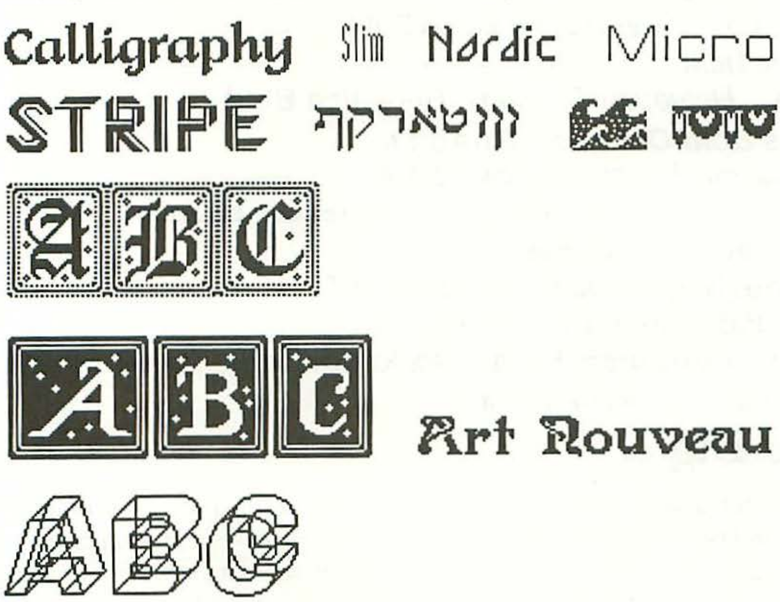

Macintosh system software uses algorithms to create italics, boldface, or other styles, unless you have installed separate fonts for them. This produces ugly results on screen and on an ImageWriter, especially for italics. You can avoid these algorithmic forms by purchasing fully designed italic and bold fonts from independent companies; Apple may distribute such fonts as well. Or, if you have patience, you can create your own italic font with a font editor. (See page 193.)

Fonts can be used not only for text but also for pictures and symbols. With a font consisting of symbols, such as Adobe's Sonata, you can display musical notes; other fonts, made up of little pictures, let you design furniture layouts with a few keystrokes. Symbol fonts can be used in any program that supports multiple fonts; they are especially valuable in graphics programs.

The Macintosh System file format sets a limit on all font sizes and allocated storage space: a maximum of 127 points and 32 KB total file size with the format used by Apple in 1988. There is not enough storage space within this limit for all the characters in fonts larger than about 60-point size, so if you create or modify such fonts, you might have to split uppercase and lowercase letters into separate files.

PostScript Fonts

When you print them on a laser printer at only 300 dots per inch, PostScript fonts are not typeset quality, although they can credibly simulate traditional typefaces.

Times Roman *Times Italic* **Times Bold**
Times Bold Italic Outline
Helvetica *Helvetica Oblique* **Helvetica Bold**
Helvetica Bold Oblique Shadow
Helvetica Narrow *Helvetica Narrow Oblique*
Courier *Courier Oblique* **Courier Bold**
Courier Bold Oblique
Symbol: αβχδεφγηιφκλμνοπθρστυϖωξψζ
Palatino **Palatino Bold** *Palatino Italic*
Bookman **Bookman Bold** *Bookman Italic*
Zapf Chancery *Zapf Chancery Bold*
variable size **Big** small

Unlike screen fonts, which consist of a set of pixels, PostScript fonts consist of a set of outlines. Italic and boldface fonts have their own outlines; this is a great improvement over the simple algorithms used for screen fonts. Each PostScript character has a single outline from which a printer's PostScript processor can create and print a character of any size. This outline is optimized for 12-point characters, however, so large characters look slightly out of proportion. Separate sets of outlines could be made, however, for large display sizes.

If you want to use a PostScript font that is not built into the printer, you will need a font outline file in your Macintosh as well as a matching screen font. The printer driver software will automatically transmit (download) the font outlines to the printer when you use that particular font. Such automatic downloading works well but takes extra time. If you know in advance that you need a particular font, you can move it manually with Apple's utility program, PostScript Dump, which is distributed by several companies, including Century Software (Berkeley, CA). Because of printer memory limitations, you may be able to download only two to four fonts at a time to the LaserWriter or LaserWriter Plus; the LaserWriter IINT, the LaserWriter IINTX, and other PostScript printers with more memory can handle a larger number of downloaded fonts.

In a PostScript LaserWriter, spacing between letters follows a simple geometric rule as type size changes. For small sizes—those less than 10 point—this rule sets the text too tightly; large sizes are set a little too loosely. Some software lets you expand the spacing between small letters for legibility.

The typesetter's method of kerning—adjusting the spacing between letters—is possible with a laser printer. Kerning is particularly useful in headlines. For example,

W A VES becomes WAVES with kerning.

LaserWriters permit kerning in 1/300-inch increments. The amount of white space between specific letter pairs (kerning pairs) is built into PostScript fonts, but kerning itself must be done from Macintosh application software. Page-makeup programs and some word-processing programs are the first ones to support kerning on the Mac.

Today, there is a growing number of PostScript fonts to choose from. Designs from Adobe (creator of the PostScript language) are first-class, but they are expensive by microcomputer standards. They can be used on one printer at a time or on five printers if you get the multiple-installation version. Adobe's fonts include two ligatures, or joined letters: *fi* and *fl*. Unfortunately, the ligatures *ff*, *ffi*, and *ffl* are missing. To handle ligatures properly, editing software should create and split such characters as required.

PostScript fonts are available from many other vendors. Century Software has the largest collection, including its MicroFonts Plus package, which produces narrower, wider, or smaller versions of fonts built into LaserWriters, and its Designer series, based on traditional typefaces. CasadyWare (Carmel, CA) also produces a line that is particularly strong in display and script designs. Casey's Page Mill (Englewood, CO) produces Bullets & Boxes, a font containing bullets, boxes, arrows, and other graphic elements. Century Software and Postcraft International (Valencia, CA) both offer tools for creating special display effects, such as shadows and elaborate outlines.

BODONI Bold *Italic* ***Bold Italic***
SANS SERIF *Italic* Book *Italic* Demi *Italic*
RIGHT BANK / RITZ *Italic* Condensed
MONTEREY *Italic* Medium **Bold** ***Bold Italic***
Calligraphy / Regency Script
Prelude Script Bold Light Slant Bold Slant
Coventry Script / Zephyr Script
Gregorian / Dorovar Italic
Кириллица *Кириллица* **Кириллица** ***Кириллица*** Russian & Ukrainian
Bodoni Ultra *Italic* Condensed Cond Italic

SANS SERIF BOLD *Italic* Condensed Cond. *Italic*
Sans Serif Extra Bold *Italic* Cond. Cond. *Italic*
GATSBY Light Demi *Italic* Demi *Italic*
MICRO Laser **Bold** *Italic* ***Bold Italic***
Micro Extend. *Italic* **Bold** ***Bold Italic***
GalileoRoman **Bold** *Italic* ***Bold Italic***
CAMPANILE / Giotto & Giotto Bold
Alexandria **Bold** *Italic* ***Bold Italic***
JOTT **Bold** *Italic* ***Bold Italic***

Fluent Laser Fonts are PostScript fonts from CasadyWare.

LPAthina Ἦμος δ' ἠριγένεια φάνη ῥοδοδάκτυλος

LPBenares यह सर्वक्षुत है कि उस समय साम्राज्यवादी

LPHebrew בְּרֵאשִׁית בָּרָא אֱלֹהִים אֵת הַשָּׁמַיִם

LaserPerfect Fonts from NeoScribe International include several non-Roman alphabets.

Screen Fonts to Match PostScript Fonts

To prepare a document for printing with the fonts built into a PostScript printer, you need screen fonts that match PostScript outline fonts. Each screen font has a built-in width table specifying how wide each character should be on the screen so that it matches printed output.

Apple supplies matching screen fonts to go with the LaserWriter and LaserWriter Plus, but these are less than ideal. They suffer from poor legibility, printed lines are too close together, and font variants such as boldface and italics are created by algorithm. Italics are particularly hard to work with because the characters lean far to the right, making it difficult to position the insertion point. Adobe produces a much better set of screen fonts that are legible and well spaced; boldface and italics come as distinct fonts.

*This screen font is Times Italic from Apple;
it was created by algorithm.*

*This screen font is Times Italic from Adobe;
it was designed as a distinct font.*

Font Editors

FONTastic, version 2.0 (Altsys, Plano, TX), is an excellent utility program that lets you manage, edit, and create screen fonts. It duplicates all the font functions of Apple's Font/DA Mover and does much more. You can copy and paste characters from font to font, change selected characters, resize fonts, and fix rough edges with a set of MacPaint-like tools. If you need special symbols, you can create or paste them into a standard font—perhaps replacing a foreign-language character you never use, for example.

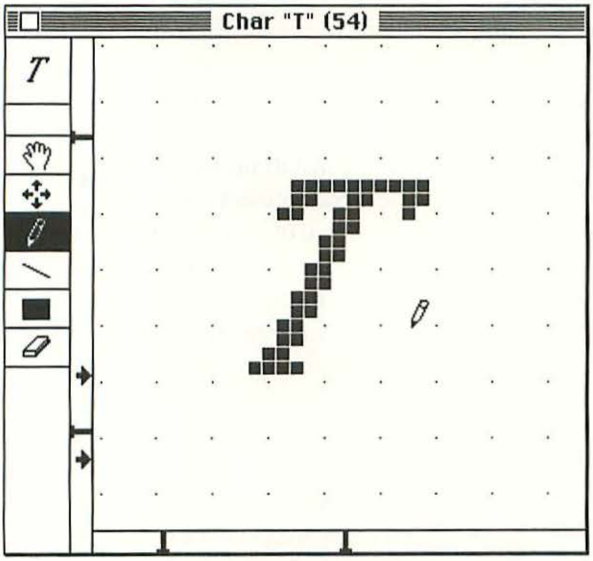

With a program like FONTastic and some patience, you can design your own fonts:

Altsys also sells Fontographer, an editor for PostScript fonts. Fontographer lets you modify an existing laser printer font or take on the ambitious task of creating a whole new one. Making new characters calls for considerable designing skill; most fonts for the many non-PostScript laser printers on the market look crude. The first good-quality laser printer fonts are now beginning to appear; the best of these are designed specifically for 300-dpi xerographic rendition. (For a good discussion on designing digital fonts, see "Digital Typography" by Charles Bigelow and Donald Day in *Scientific American,* August 1983.)

Managing Fonts

Although you can install many fonts on a Macintosh, you might not be able to use every font with every program. Application programs that follow Apple's guidelines can handle as many fonts as you can install in the System file plus the application. Applications that do not follow the guidelines (such as early versions of MacDraw) may give you access to only a limited number of fonts.

Many applications restrict the number of sizes you can use on screen or limit you to a few fixed sizes; they do no scaling. If you print on an ImageWriter or LaserWriter IISC, however, such restrictions do not matter because you must print from installed fonts anyway. The most flexible programs, which let you work on screen in any font size you want, installed or not, complement PostScript printers, which can create a good-looking font of any size for printing.

With Apple's early 1988 system software, you can install up to 256 screen fonts on one disk, but the practical limit is about 200. (The desk accessory Suitcase lets you get around this limitation; see Chapter 14.) Because each size of each typeface is a separate font, it occupies a separate slot. Each screen font is assigned a number so that the system can keep track of them all. For a single-user system, it usually doesn't matter which number goes with a particular font as long as the installed fonts have distinct numbers; Font/DA Mover reassigns font numbers when necessary to prevent confusion.

The numbers assigned to fonts do matter, however, if you are on a local area network or work regularly on more than one Macintosh. Unless the same font files have been installed on each machine, different fonts could have the same numbers, or the same font could have different numbers, producing some unexpected results when you print.

The font-number limitations and potential mix-ups should no longer be problems when new Macintosh system software is able to put font information into a separate font file containing identifying information such as font name, point size, and manufacturer.

Typographic Tips

With so many fonts easily accessible, many people are tempted to use them all—on a single page. But a page that looks like a kidnapper's ransom note is unattractive and unreadable. Well-designed books and magazines actually use very few fonts; you'll rarely see more than two typefaces on a page. Too much boldface or too many italics ultimately diminish the emphasis they are designed to give. And embellishments like outline or shadow belong only in display type, not in ordinary text. (For an excellent survey of typography and aesthetics, see *The TypEncyclopedia* by Frank J. Romano, R. R. Bowker, 1984.)

Many Macintosh fonts include symbols, such as •, ¶, ®, ©, and true open and close quotation marks (" and ") not found on typewriters or most other computers. Most fonts also have three dash lengths: hyphen (-), en-dash (–), and em-dash (—), each of which has its special grammatical function in print. (See *The Chicago Manual of Style,* 13th edition, University of Chicago Press, 1982, for how to use these marks.) Especially with the PostScript fonts available on the LaserWriter, these marks look more professional than simple quotes (' ") or a double hyphen (--) as a dash. Caution: Do not use these special Macintosh characters in any file you intend to transfer to another computer that cannot recognize them.

Although most books and magazines are printed with fully justified (flush) left and right margins, studies of legibility have generally shown that text with a ragged right margin is easier to read. Ragged left is even more legible than ragged right, but ragged left is rarely used except for display purposes.

To produce attractive right-justified text, particularly on a page with two or three columns, you need to hyphenate words; otherwise, the computer simply inserts extra space between words, and your text looks loose and is difficult to read. If you plan to right-justify your text, use an editing program that permits conditional, or discretionary, hyphenation so that words break only when necessary.

Manual hyphenation is tedious, so many companies have developed automatic hyphenation programs. These programs insert hyphens by following a set of rules or a word list. Hyphenation by rule breaks words according to conventions that most dictionaries follow. A small exception dictionary covers common words that do not conform to these rules. Hyphenation by word list breaks words only where the list indicates; words not on the list must be hyphenated manually. Depending on the text and the size of the list, you might still end up manually hyphenating a lot of words.

No hyphenation program can be completely automatic. The noun *pres-ent* breaks differently from the verb *pre-sent*. The best hyphenation programs flag such homographs; they also warn you if there will be more than three hyphens in a row (which looks awkward) and include a separate set of rules or a word list to adjust for variations between American and British hyphenation.

GRAPHICS AND GRAYS

In addition to typeset text, publications usually contain line art (such as cartoons, diagrams, and graphs) or pictures (such as drawings, paintings, and photographs). Whether it is done on a computer screen or on the printed page, reproducing type and line art requires only black and white, but realistic pictures need grays. Grays can be produced several ways, depending on both hardware and software.

True grays, which can be produced by video screens or photographic processes, look gray even under magnification. They come in two forms:

- Continuous grays, consisting of every intermediate tone between black and white. Among common image-reproduction processes, only photographic films and prints and analog video systems can show continuous grays; printing presses cannot print them.
- Stepped grays, consisting of only specific shades of gray, which are most often produced by digital techniques in discrete steps. The greater the number of steps, the closer an image built from stepped grays can be to one done with continuous grays. (Color printing presses can print stepped grays if they are set up with a series of gray rather than color inks.)

Dithered grays are produced by devices that can display only black dots of fixed size. The grays are simulated by scattering black dots within an image so that certain regions appear gray. The dots can occur repetitively in regular numbers and patterns, each corresponding to a distinct shade of gray, or they can be scattered in a pseudorandom pattern, in which the density of dots determines the shade. Repetitive patterns make it easier to distinguish particular tones—for example, in images where the distinction among grays, not shapes, is important. Pseudorandom patterns usually give more realistic images. In images with dithered grays, the black pixels serve double duty: Their position defines an object's contours, and their pattern or density defines its shading; the rendition is necessarily coarse.

Dithering can also be applied to stepped gray images to create the appearance of even more shades.

These photographs show the difference between a repetitive pattern of dithered grays (left, from the MAGIC digitizer) and a pseudorandom pattern (right, from the MacVision digitizer). Courtesy MacWorld *magazine.*

Halftones are grays built up from evenly spaced dots that vary in size; the larger the dots, the darker the gray. Photographic prints are made from halftones in which dot size varies on a continuous scale. Laser printers or phototypesetting machines can produce only stepped halftones, which are made of variably sized dots that are in turn made up of smaller dots of fixed size. In contrast to dithered grays, halftone image sharpness is determined by the number of dots per inch (or lines per inch), which is constant throughout a halftone image; contours are defined by dot position. The halftones in newspapers are printed at 65 to 85 lines per inch (lpi); those in magazines are typically printed at 120, 133, or 150 lpi. High-quality art reproductions may be printed at 200 to 300 lpi.

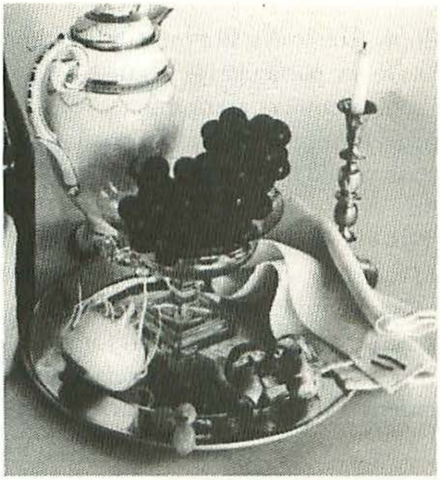
A


B


C


D

©*MacWorld*. January 1988 (5:1)

Picture A is reproduced by conventional photographic halftone techniques at 100 lines per inch. Picture B was scanned into a Mac by an Abaton scanner with 16 levels of gray and printed at 100 lines per inch on a Linotronic 300 typesetter. Picture C was scanned by a Datacopy scanner with 64 levels of gray and printed at 100 lpi on a Linotronic 300 typesetter. Picture D is the same as picture B except that it was printed at the equivalent of 50 lpi on a LaserWriter. Courtesy Macworld *magazine.*

Grays on the Macintosh screen

The screen in the small Macintoshes cannot show true grays, but only dithered grays made from black and white pixels. These grays are not particularly effective in displaying the rich image that can be generated by some scanners and printers. For a more detailed rendering, you could magnify the image, thus providing more pixels to work with; but the screen is so small that magnification would mean that only a tiny fraction of the overall image could be visible at one time.

Video cards for the Mac II store multiple bits for each pixel and can support true stepped grays. Apple's Macintosh II Video Card can store 4 bits per pixel, or 16 levels of stepped grays; if you combine it with Apple's Video Card Expansion Kit, you can get 8 bits per pixel, or 256 levels of stepped grays, which still aren't enough for smoothly toned images. Video boards that display 24 bits per pixel, or 1,677,216 levels of stepped grays, are now available. Of course, Mac II screens can also display dithered grays or dithered stepped grays.

Grays printed by the ImageWriters

The ImageWriters can print only dithered grays. The ImageWriter LQ does this best because it prints the smallest dots at the highest density, but no ImageWriter is really suitable for images that need a full range of grays. (Another way to get these printers to print grays would be to use a ribbon with four gray tones instead of colors; stepped grays would then be possible, and you could create the image by means of color drawing techniques, except that each color would be a step of gray. No one has yet produced such a ribbon, however.)

Grays printed by the LaserWriter

Laser printers, like ordinary dot-matrix printers, can print only black. But at 300 dpi, they store enough pixels to print stepped halftones at 60 lines per inch in both dimensions, which gives a recognizable image—but one of lower quality than a newspaper's. At 60 lpi, each halftone dot can contain up to 25 LaserWriter pixels—5 horizontally and 5 vertically. Each black pixel added to a halftone dot produces a slightly darker shade of gray; 25 black pixels print as black. In principle, you could print a halftone image at any resolution, but the finer the resolution (the more lines per inch in the printed image), the fewer the pixels available for each halftone dot and, therefore, the fewer shades of gray you would have. If you printed at 75 lines per inch to get a sharper image, for example, you would have only 16 pixels per halftone dot—4 in each dimension—and only 16 shades of gray.

Producing a halftone on a LaserWriter is slow, requiring many minutes of processing.

Grays printed by PostScript typesetters

Like laser printers, PostScript (and other) typesetting machines can create only stepped halftones, but their pixels are so much smaller that much finer tonal gradation is possible. However, at 120 lines per inch—the resolution of the halftones printed on most magazine pages—the steps are still visible. For this reason, magazine halftones are produced by photographic, not phototypesetting, means. And, as with a laser printer, creating halftones on a typesetter under PostScript control is slow.

SCANNERS AND VIDEO DIGITIZERS

The power of Macintosh graphics programs has almost obscured the fact that creating attractive images takes both skill and an aesthetic sense. Professional artists might be able to create satisfying pictures directly on the Mac, but the rest of us need help. The quickest and easiest way to get a picture into the Mac is to use a scanner, an input device that converts existing artwork or a video signal into a set of pixels for processing by the computer. Once it is in Macintosh file format, the image can be altered or combined with text or pictures from other Mac files, and the result can be printed. One caution, however: Because a scanner can function essentially as a copier, copyright laws may apply to images you scan.

Scanners come in several forms for several purposes. Page scanners contain image-sensing hardware plus digitizing circuitry for converting a printed image into bits; a video digitizer contains only digitizing circuitry, relying on the output of a separate image sensor, such as a video camera, to scan the object or image. Page scanners can be further divided into array scanners and printer scanners; how all three scanner types operate and the kinds of information they store vary widely.

The best scanners are modestly priced and soon pay back their cost in the time they save. Scanners can be connected to a Macintosh through the SCSI port or a serial port; most now connect via the SCSI port because its information transfer rate is much higher.

Array Scanners

Array scanners capture images—either graphics or text—from a paper original by means of a linear array of photosensors that move across the page, sensing changes in light values within the image. The scanner circuitry monitors photosensor output and then sends it, line by line, to the computer. Many scanners sold for computers are simply modified facsimile (fax) machines, which send images over standard phone lines.

The resolution of the image a scanner sends to the computer is determined by the number of photosensors in the array as well as the rate at which the scanner samples photosensor output. The most common scanners contain 2555 photosensors, which give a resolution of 300 dots per inch across an 8½-inch page. Most scanners can operate at several resolutions to generate images that can be processed by a variety of

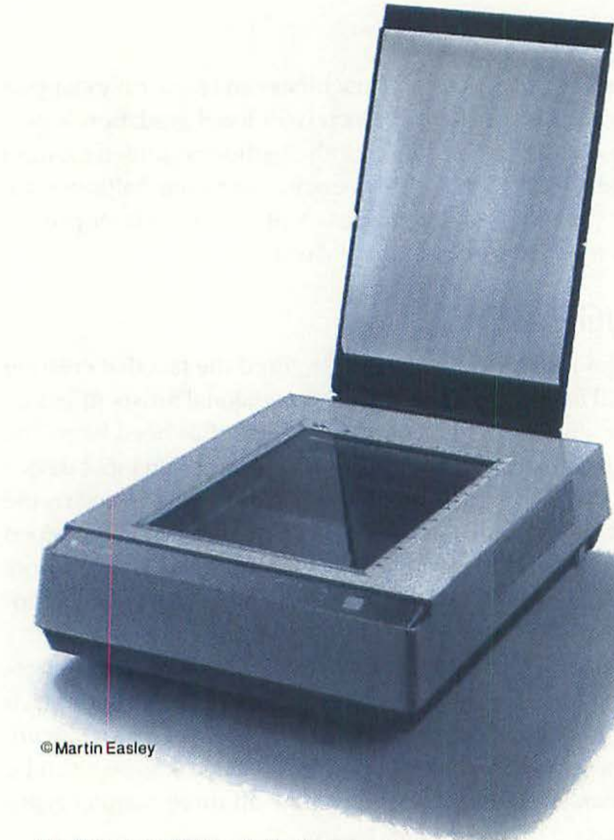

The Datacopy 730 is a flatbed array scanner.

devices, not only computers. At 300 dpi, for example, the scanned picture can be printed directly by a laser printer (provided the laser printer has full graphics capability); at 200 dpi, the scanned image can be sent via Group III facsimile (see Chapters 13 and 19); and at 75 dpi, the image can be displayed on a computer screen or printed by low-cost dot-matrix printers.

The optics in an array scanner are fixed with respect to the paper; no zoom-lens arrangement is possible. Therefore, for a scanner to operate at several resolutions, the 300-dpi array must be altered in some way. To create a 200-dpi image, the scanner ignores the output of every third photosensor; to create 75 dpi, it can either ignore three of every four photosensors or average the output of four photosensors. To modify the resolution along the dimension perpendicular to the array (in the direction of motion), the scanner changes its sampling rate; 300 samples during one inch of motion yield 300 dpi along the page.

The simplest, smallest array scanners can only scan images from loose sheets of paper; the photosensor array is fixed, and the originals move past it. Larger, more expensive flatbed scanners work like photocopiers: You lay down anything you want to scan, including bound books, on a glass plate, and the array moves beneath the stationary original. Flatbed scanners are better for most jobs than sheet-fed scanners, whose paper-moving mechanism can mangle originals.

Array scanners from different manufacturers—most are made in Asia—work similarly and offer similar features. Some are built better than others, but as long as the scanning mechanism is properly aligned, there is little hardware difference among models.

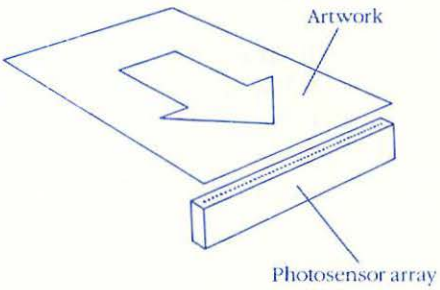

In array scanners, photosensors span the width of the paper. In some models, the sensor array moves down the paper; in others, the paper moves past the sensor.

Printer Scanners

These scanners take advantage of a scanning platform many Mac owners already have—the ImageWriter or ImageWriter II. In the low-cost ThunderScan system (Thunderware, Orinda, CA), the standard printer ribbon cartridge is replaced by a scanner cartridge containing a light source and a photoelectric sensor. You roll the artwork into the printer, and the ThunderScan software runs the printing mechanism, moving the artwork up while the sensor scans horizontally as if it were printing. ThunderScan can digitize an area up to 8 by 10 inches. Epson (Torrance, CA) produces a similar scanner for some of its printer models.

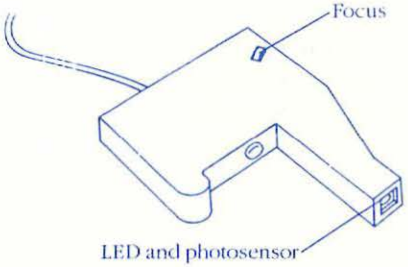

The ThunderScan cartridge replaces the ribbon cartridge on an ImageWriter.

Video Digitizers

Instead of scanning a printed original, video digitizers convert the output of a television camera or other video source—including broadcast programs, videotape, or some computers—into pixels for computer processing. With a video digitizer you could, for example, digitize a picture taken by a home video camera and transmit it to your Mac. (Although video digitizers can convert the display output of some inexpensive home computers, they cannot deal with the signal generated by most computer displays; to be convertible, a signal must conform to the RS-170 standard.)

Prices of video digitizers, excluding a camera, range from about $200 to $600; units with special features can cost as much as $2,500. One low-cost monitoring device, Micron Technology's MicronEye Optic RAM (Boise, ID) contains a tiny video digitizer and a built-in camera. Unlike a traditional video camera, the device picks up a video signal by means of a rectangular solid-state array that is actually a RAM chip without a cover; the chip's memory cells are light sensitive.

Resolution

Scanner resolution is most often quoted as the total number of pixels a scanner can generate for a particular image size, such as an 8½-by-11-inch sheet of paper or a single video frame. Most array scanners and the ThunderScan are designed to scan a sheet of paper, usually at full size. The following list gives the resolution (in this case, the same as pixel count) of some common scanners.

Array scanner	
200 dots per inch	1728 by 2200
300 dots per inch	2400 by 3300
ThunderScan	
1.0X magnification (72 dots per inch)	576 by 720
1.12X magnification with gray scale (80 dots per inch)	645 by 806
2.76X magnification without gray scale (200 dots per inch)	1589 by 1987
Video digitizers	
Koala MacVision	640 by 480

The resolution of the final image—once it has been scanned, manipulated in the computer, and printed again—depends not on the scanner's resolution but on the printer's. If you scan an image at a higher resolution than you can print, you won't be able to get the image out of your printer at the higher resolution without magnification. It is therefore best to scan at or near the same resolution your printer uses.

The output resolution of a video digitizer is limited by the quality of the incoming video signal. Few, if any, television cameras have sufficient resolution to faithfully reproduce intricate detail, such as black and white pixels arranged alternately along a horizontal scan line. Nevertheless, a digitizer that can generate a large number of pixels to send to the computer can position images more precisely and render curves more smoothly than one that generates few pixels.

Color in a video signal can also limit image quality. Because color information obscures fine details, some digitizers include a filter that can be turned on or off to remove the signal's color component. The best cameras to use with a video digitizer are high-quality black-and-white models designed for precision closed-circuit TV.

Maximum vertical resolution of video digitizer output is determined by the total number of lines in the incoming video signal. The television format (NTSC) used in North America and Japan imposes a limit of 484 lines per frame (NTSC format is often cited as having 525 lines per frame; this figure includes the 41 black lines in the black bar that separates frames, called the blanking interval.) Digitizers that can handle the European CCIR format (the basis of PAL and SECAM color standards) can get up to 575 lines (usually quoted as 625 lines including the blanking interval).

All television systems produce interlaced images, in which a single frame is broken into two vertically offset fields—one of 242 even-numbered lines (287 in CCIR) and one of 242 odd-numbered lines. The odd-numbered and even-numbered fields are supposed to intermesh exactly to give a complete picture, but many low-cost television cameras cannot interlace precisely; often one field is simply superimposed on the other. If so, you cannot get more than 242 lines of usable vertical resolution, regardless of the digitizer's capabilities.

Because video images inherently have very poor resolution, the black-and-white output of a simple video digitizer is usually unsuitable for printing by much higher-resolution output devices such as laser printers and typesetters. If, however, the image contains gray-scale information, a laser printer or typesetter can use this information to produce halftones.

Geometric Accuracy

Will a straight line in an original document be a straight line in a scanned image? It depends on the scanning system. Array scanners are highly accurate. The ThunderScan is limited by the precision of the ImageWriter mechanism. Getting accurate vertical lines is the most common ThunderScan difficulty; because of imprecise print-head positioning, successive passes may be shifted one pixel to the left or right. If this happens, try turning the image sideways in the printer.

Video digitizers are subject to many small defects common in the video cameras whose signals they digitize. Video camera lenses, particularly zoom lenses, often exhibit visible barrel or pincushion distortion. Further, the camera pickup tube (based on

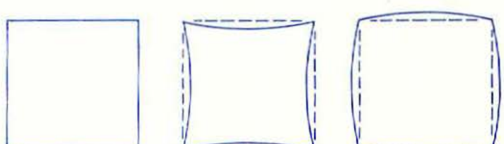

Pincushion and barrel distortion.

vacuum-tube technology), which converts the image projected by the lens into an electronic signal, suffers from nonlinearities that get worse with age. Professional video equipment incorporates circuitry to compensate, but the costs are too high to include such circuitry in simpler equipment. The newer solid-state cameras and the MicronEye have perfect image-sensor geometry, so they are limited only by the accuracy of their lenses.

Scanner Speed

The speed of a scanner or digitizer can be considered as having two distinct steps: the length of time it takes to capture an image and the length of time it takes to digitize it. For page scanners, which scan more slowly than video digitizers and nearly always work on stationary images, these two steps are the same. Array scanners take from 5 to 20 seconds to scan a full page; the ThunderScan takes 10 to 20 minutes, or about 5 minutes for a pixel count equivalent to a video frame.

For most video digitizers, image capture and digitizing time are also equivalent—typically 4 to 30 seconds. Like page scanners, such units are suitable only for stationary objects or images that have been frozen by other devices, such as videocassette recorders (VCRs). Many of today's VCRs and televisions have a built-in frame buffer (image memory) for freezing action; without such a buffer, the image produced by the still-frame function of most home VCRs is usually too unstable to produce a distinct digitized image (although videodisc players work fine).

To deal with moving images, some video digitizers themselves contain a frame buffer, which greatly shortens capture time; once the image has been captured, digitizing can be done at leisure. Frame buffers normally freeze a frame in 1/30 second; some can freeze a single vertical field in 1/60 second. Specialized video cameras can freeze faster action, and their output can be digitized via a frame buffer.

If you are trying to take a video photo of an object you want to digitize and are using the digitizer's output to focus the camera, a slow digitizer can be annoying. It's better to plug in an ordinary video monitor to get a live image for focusing and framing.

Optical Character Recognition

One widespread use for scanners that has nothing to do with graphics is optical character recognition (OCR), in which typed or printed text is read directly into the computer without your having to retype it. OCR software examines a character's bit map, compares it to templates stored in memory, and identifies the character, which is then passed to the computer as text, not graphics. The software can be an integral part of the OCR machine itself, or it can run as a microcomputer application.

Although simple in principle, OCR is difficult in practice. Since even a one percent error rate during scanning means about 15 errors on a page, recognition accuracy must be very high. Most OCRs are fussy, able to read only a few monospaced typewriter-style fonts, such as Courier; most cannot read proportionally spaced type, although some flexible OCR systems can be "trained" to analyze new fonts. The best OCR systems are still expensive ($15,000 to $40,000) and incorporate their own processing hardware and software.

File Formats

In the early days of the Macintosh, a scanner or digitizer only needed to save files on disk in MacPaint format. Now, with many more graphics programs available, there is a profusion of file types; good supporting software for scanners ought to generate most or all of them. Virtually all graphics application programs for the Mac can read at least one of these formats:

- MacPaint: bit-mapped format limited to a single page at 72 dpi; only dithered grays possible.
- PICT: a format developed by Apple for exchanging attribute-based drawings among graphics programs. PICT is best for describing objects, such as lines; it can also describe bit maps, which is how scanners use it. PICT can deal only with dithered grays; the revised version, PICT2, adds stepped grays, colors, and other features.
- Tag image file format (TIFF): developed by Aldus and Microsoft; includes stepped grays. TIFF is not completely standardized; many incompatible variants are in use.
- Encapsulated PostScript format (EPSF): saves the image as a set of PostScript commands that can be sent to a PostScript printer along with other PostScript commands; can handle stepped halftones. This is an excellent format for mixing pictures with words or numbers in the same printed document. For example, the scanned EPSF file can be transferred to a database program, which need not "understand" it; the database program can simply pass the EPSF file on to the printer along with its own PostScript commands. Unfortunately, EPSF files generated by different graphics programs are sometimes incompatible.
- LaserBits: bit-mapped format for 300-dpi laser printer images developed by Silicon Beach Software (San Diego, CA) for SuperPaint.
- Group III fax format: for facsimile transmission. (See Chapter 19.) Dithered grays possible.
- Optical character recognition (OCR) output: best generated by 300-dpi scanners for creating a text file.
- Native, or proprietary: a file format used only by a specific scanner manufacturer. A native format is sometimes needed to support special features or to compress a file into a more manageable size.

Whatever their format, scanner output files are usually large. MacPaint format is the most compact. It stores a 72-dpi image with one bit per pixel and uses compression—that is, a long string of black or white pixels is coded into a single number rather than as separate pixels. A typical scanned page might occupy 50 KB of memory. The image generated by a 300-dpi scanner easily consumes half a megabyte or more without grays and two to five megabytes with grays. High-capacity hard disk drives are essential with these scanners.

Working with Scanners

Not all the images you will want to scan into your Mac will be the right size. If you're using a page scanner, resizing is best done with a photostatic camera or an enlarging-reducing photocopier before scanning. If you're using a video scanner, resizing can be done on the spot with a zoom lens or by moving the camera. You can also resize an image after scanning with a paint program that works with bit maps, but the scaling process introduces distortion.

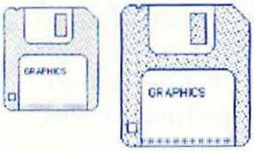

Stretching an image in MacPaint distorts some objects.

After scanning, you might want to edit or redraw certain images before placing them in a document or printing them. The available graphics programs vary as to the kind of scanned images you can manipulate and how. Some programs can handle gray-scale information passed to the Mac from the scanner. Some allow you to edit images bit by bit, whereas others let you translate bits into lines, curves, or polygons for editing. Some generate output suitable for printing only on a LaserWriter, and others only on an ImageWriter.

The sensors in many scanners detect grays. These scanners usually give you some sort of contrast adjustment so that you can arrange the distribution of gray steps between black and white. Set for high contrast, the scanner should drop grays altogether and produce a purely black-and-white image. Once the grays have been sent to the Mac, the graphics program ImageStudio (Letraset, Paramus, NJ) lets you adjust them.

Simple scanners do not send gray-scale information to the Mac; they simply dither the grays and then discard the original gray information. The only way to change the shading on such an image is to rescan it.

To enable you to edit 300-dpi images, scanned or drawn, on the Mac's 72-dpi screen at full resolution, SuperPaint has a feature called LaserBits that magnifies the image four times; at that magnification, one 72-dpi dot corresponds to one 300-dpi dot when the picture is printed at its original size. Several other painting programs, including Canvas (Deneba Software, Miami, FL) now include this feature. But be careful: At 72 dpi, you can clean up your work pixel by pixel; with 300 or more dots per inch to work with, doing so could take forever.

The converse strategy enables you to edit at an effective 300-dpi resolution with a standard bit-oriented paint program: Edit the image on the screen and then move it via the Clipboard to an object-oriented graphics program, such as MacDraw. From there, print the image at 25 percent of its original size.

All scanners create bit-mapped images. Yet for a good deal of graphics work, an image defined in terms of objects is much easier to manipulate. Consequently, much effort has gone into developing object-recognition software. Object recognition tries to do for the components of a drawing what optical character recognition does for text: translate groups of pixels into codes for lines, circles, and angles. This goal has proven formidable.

So far, most object-recognition programs can detect only lines. But because what you intend to be a single line may cross many other lines, these programs tend to detect more lines than you want in intricate drawings; a cross, for example, might be detected as four separate lines, and each piece might have to be manipulated separately. In other cases, object-recognition programs might detect fewer lines than you want; they cannot, for example, find hidden lines in a three-dimensional perspective drawing where the front of the object obscures the back.

Digital Darkroom (pre-release, Silicon Beach Software) gets one step closer to object recognition by recognizing contours. It takes scanned images containing grays and constructs contour lines by connecting pixels with the same gray-scale value.

Until object recognition improves, programs like Adobe Illustrator and Aldus FreeHand can help fill in. They let you manually construct an object-oriented picture on top of the bit map of a scanned image. Using special tools for different kinds of objects (a French curve, for example), you specify the end points or limits of an object, and the program stores it; when you are done, you can discard the original bit-mapped image.

PAGE MAKEUP

Pasteup is the process in which a graphic designer physically lays out and pastes down all the components of a page. In addition to text and pictures, a page can contain rules (lines that separate page components), decorative borders, or initials (sunken or decorative initial capital letters used to set off a portion of text). Text can be divided into several columns; headlines might appear anywhere on the page; pictures might have text wrapped around them. Placing these elements on a page and rearranging them electronically is what desktop publishing is all about.

Page-Makeup Programs

Because page-makeup programs are changing rapidly, this short survey cannot hope to cover the field. For more—and more recent—information, see magazines and the newsletter *Desktop Publishing,* by Tony Bove and Cheryl Rhodes (PCW Communications, San Francisco, CA).

PageMaker (Aldus, Seattle, WA), the leader in the page-makeup market, enjoys the most support from training classes, books, and software templates. The program, whose tools are easy for a graphic artist to master, is well suited to publications where many pages have different designs. Version 3.0 (pre-release) includes style sheets for specifying the layout and format of a particular document. It supports scanned images better than other page-makeup programs. One potentially serious drawback is that the maximum document size is 128 pages. So far, PageMaker is the only page-makeup program that can not only read word-processing files but also generate them (in Microsoft Word and WriteNow formats).

ReadySetGo, version 4.0 (Letraset, Paramus, NJ), also has style sheets but no restriction on the number of pages, so it can be used for books as well as for shorter publications. It has basic word-processing features; good hyphenation; and the ability to search for text according to font, style, and type size.

Although every page-makeup program described here can edit text, most assume that you will create and edit long documents in a separate word processor. XPress, version 1.1 (Quark, Denver, CO), is unusual because it contains a full word processor inside a page-makeup program. XPress has as many composition features as PageMaker and ReadySetGo, but it lacks style sheets.

Scoop, version 1.0 (Target Software, Miami, FL), concentrates on graphics tools. It contains a word processor plus a paint and draw program, so you do not have to leave the program to create or retouch graphics. It does not have style sheets, however, and publications are limited to 100 pages.

MacPublisher III, version 3.01 (Boston Publishing Systems, Boston, MA), contains some features valuable to graphic designers, such as the ability to rotate text and pictures, but overall it has fewer features than the other general-purpose page-makeup programs described here. It is also the least expensive.

Ragtime, version 1.1 (Orange Micro, Anaheim, CA), is specifically aimed at producing financial publications and other documents with numeric tables. It lacks many typesetting features found in other programs, but it does contain a spreadsheet. Cells in the spreadsheet can be individually formatted for fonts and styles, and spreadsheet data can be linked to publications.

Interleaf Publisher, version 1.0 (Interleaf, Cambridge, MA), is the most comprehensive page-makeup program to appear so far for the Macintosh. Designed for large technical documents, it contains many computer-aided design features, including three-dimensional perspective diagrams and a full word processor, but it lacks the kerning and spacing controls that are standard with other page-makeup software. The program was originally written to run under the UNIX operating system for Sun, Apollo, and DEC workstations. The Mac version, which uses the Mac operating system, offers the same functions as its workstation counterparts, except it does not support scanned graphics and will drive only PostScript printers. Interleaf Publisher can use only its own fonts, not those installed in the System file. All versions, workstation or Mac, generate compatible files.

Despite its modern design, Mac users may find operating Interleaf Publisher a little awkward at first because the company has retained the windowing interface that it developed for its earlier software; the program does not have a Macintosh interface except for some basic file maneuvers.

Interleaf Publisher is big, requiring a Mac II with 5 MB of RAM; the program files alone take up 8 MB of hard disk storage. It is also expensive—$2,500 compared with the $100–$700 of the other programs described above.

FINISHED COPIES

For the cleanest results with an ImageWriter, print your pages oversized and reduce the image—with a reducing photocopier or photographic process—for the final copy; the reduction decreases the effective dot size and increases the visual density of the dots. Even LaserWriter output can benefit from reduction. The advantages disappear, however, if you must paste up pages again—the very process that all this software and hardware is supposed to replace.

Typographic Composition

Microcomputers have been used for several years to prepare text for typesetting equipment. Special typesetting commands specifying fonts, paragraph formats, headings, and so on are embedded in the text file with a word processor; the embedded commands constitute a markup language. The process is tedious and awkward, and with all the embedded commands, the text itself is barely readable. Moreover, the coding process is not interactive; you cannot see what your page will look like until it comes back from the typesetter.

Some widely used markup languages include n-roff and t-roff from the UNIX operating system; the TeX typesetting language, developed by Donald Knuth and optimized for equations; and International Standard ISO 8879, also called Standard Generalized Markup Language. As part of its Electronic Manuscript Project, the Association of American Publishers has proposed that ISO 8879 become a publishing-industry standard.

If necessary, you can use a Macintosh to prepare text for typesetting in this traditional manner. You can also use software that works from commands embedded in your text file and generates PostScript code suitable for a PostScript LaserWriter or typesetting machine. Such software includes two versions of TeX (one from Addison-Wesley in Reading, MA, and the other from FTL Systems in Toronto, Ontario) and Just-Text (from Knowledge Engineering, New York, NY). None of these programs are interactive, although some of the newest versions might let you preview the finished page on the screen.

But although these markup languages work, they are the antithesis of the Macintosh. Why live with embedded commands when you can create completed pages on the Mac and send them directly to a PostScript typesetter?

As desktop publishing on the Macintosh becomes more widespread, more typesetting companies are getting equipment that can handle PostScript code. Once you have created a printable document with the fonts and graphics of your choice, you can simply take a disk to one of these typesetting firms and get back clean, typeset pages. The best way to ensure that you get back what you put in is to include your Mac System file, with the necessary fonts, on the disk that goes to the typesetter. If you use downloaded fonts, be sure the typesetter has identical ones available.

PostScript convenience and precision is not restricted to text, of course. Any graphics you have created using a program that generates PostScript code can also be printed by a PostScript typesetter. Moreover, if your graphics program is not sophisticated enough to execute your most intricate work, you can send PostScript code specifying the most detailed portions of your image directly to the typesetter—something like embedding typesetting codes in text; the codes will be interpreted and executed during printing.

Typesetters

For really clean finished copies, only typeset quality will do. The Linotype Company (Hauppauge, NY) produces three typesetting systems. All three include a raster image processor (RIP), built in or external, which receives information in PostScript codes via AppleTalk and converts it into a series of lines (raster). The RIP has its own 68020 CPU and a hard disk drive for storing font information. The internal RIPs have 2 MB of RAM and a 20 MB hard disk, which can hold about 125 PostScript fonts; the external RIPs have 6 MB of RAM and an 80 MB hard disk, which can hold about 500 fonts. All models can produce pages at several different resolutions. Each model has a fixed maximum printing width; printing depths vary, depending on printing resolution.

- The Linotronic 100 (priced at $32,000 and higher, including the RIP) creates finished copy at a maximum resolution of 1270 dots per inch, with a maximum printing width of 11.7 inches and a maximum page depth of 103 inches at 317 dpi by 635 dpi. The smallest and lowest-priced configuration has a built-in RIP with a 20 MB hard disk and 3 MB of RAM; a model with an external RIP goes for $45,000.
- The Linotronic 300 (priced at $71,000 and higher, including an external RIP) achieves a maximum resolution of 2540 dots per inch, with a maximum printing width of 12 inches and a maximum page depth of 25.8 inches at 635 dpi by 1270 dpi.
- The Linotronic 500 (priced at $107,000 and higher, including an external RIP) can print at a maximum resolution of 1690 dots per inch, with a maximum printing width of 17.5 inches and a maximum page depth of 77.2 inches at 423 dpi by 846 dpi.

All three units contain a helium-neon laser that writes directly onto photosensitive paper, film, or printing plates, which must be handled in lighttight cassettes and which require photographic processing. The results are clean, with smooth, even black areas free of the stray toner that mars the white areas of some laser printer output. The dot size at maximum resolution is under one mil.

Laser typesetters achieve much higher resolution than laser printers because their dots are not smeared during the toning and fusing steps of a xerographic process. And unlike laser printers, laser typesetters are precise enough to pause in the middle of printing a page to receive more information. Consequently, they have no need for a large memory buffer to store an entire page as a laser printer does. A good thing, too—at 2540 dpi, an $8^1/_2$-by-11-inch page contains 603 million pixels, or 75 MB of data.

13: A Communications Primer

Someday, fast and efficient communication between computers will give us instant access to one another, the world's libraries, and the latest political gossip. But today, communication can be the most frustrating procedure that you attempt with your computer.

When you use the Macintosh to communicate with other computers, you enter a chaotic world with no industry-wide standards. Unfortunately, with the number of incompatible products increasing all the time, the situation is likely to get worse before it gets better. Nonetheless, with a telephone line, the right hardware and software, and much patience, you can connect your Macintosh with public databases, electronic mailboxes, facsimile machines, and other people who also have the proper equipment.

WHAT YOU NEED

To use your Macintosh to talk to another computer, you need:

- A telephone line that is not a party line.
- A modem (from *mo*dulator/*dem*odulator). Modems transmit data by converting a computer's digital signals into modulated audio tones that can travel on a telephone line. They also do the reverse: demodulate the audio tones from a distant computer and modem back into digital form. Many different modems will work with the Macintosh.
- Communication software. The software must be compatible with your modem.

Modems for microcomputers generally operate at three speeds: 300, 1200, and 2400 bits per second (bps). (The term *baud* is commonly and inaccurately used as if it were interchangeable with bits per second. For a precise definition and a table of true baud rates, see Chapter 19.)

In terms of transmission speed, 300-bps modems are slow, sending 25 to 30 characters per second (cps). At a rate of 30 cps, it takes about a minute to fill a Mac screen with characters, or up to 12 minutes to fill it with graphics. The 300-bps modems have been popular among hobbyists because of their low price—from $40 to $125. Higher-speed modems include more expensive and more precise components; 1200-bps models cost from $100 to $600. But unless you are on a very tight budget, 1200-bps modems are worth the extra money. They transmit four times as fast as 300-bps modems, at 100 to 120 cps. Farther up the scale, 2400-bps modems are twice as fast as 1200-bps models, and many companies have introduced models that run at 9600 bps or faster.

In addition to a correctly wired phone line and a modem, you will also need some information about the computer you want to communicate with and its modem. These communications parameters include:

- Speed of the other (remote) modem.
- Character width: how many bits make up each of the characters being transmitted (seven or eight); also called data bits.
- Stop bits: how many bits (one or two) mark the end of each character, so that the computer knows when one character stops and another begins; this parameter is often unnecessary.
- Parity: a simple form of error detection (even, odd, or no parity).
- Handshake: a way for the receiving computer to tell the sending computer to pause (XON/XOFF, clear to send, or no handshake); this parameter is not always needed for short messages or files.

The exact values of these parameters are not important by themselves; what's crucial is that the computers at both ends of a communication link are set to the same parameters. If you are linking up with an electronic mail service or database, you must set the parameters to conform with the remote computer. But if you are hooking up with a computer whose operator you can talk to, both of you can simply agree to use the same values of these parameters.

MAKING THE CONNECTION

Connect your modem to the Macintosh and to a telephone jack according to the instructions in the modem manual. Whenever possible, the modem should connect directly to the telephone line through a modular jack. If your phone uses a four-pronged plug, you should buy a high-quality adapter plug; cheap adapters are not reliable.

Do not connect your modem to a telephone party line. If someone on the party line picked up the phone during a computer communication, the link would be disturbed; moreover, in an emergency the other party could not break in to ask for the line.

Most modems now include automatic dialing and answering, so you can dial telephone numbers stored on disk (or entered from the keyboard), and your modem can automatically answer incoming calls. Take care to turn off your modem's autoanswering feature unless the calls coming in on that line are exclusively from computers; otherwise, a caller will be greeted with an irritating, high-pitched tone.

Long-Distance Communications

The efficiency of modem communication varies with phone line quality. Modem connections on local calls generally work satisfactorily, but long-distance calls frequently run into problems because of noisy lines. You may find that some long-distance services have noisier lines than others. The line quality can also vary with the time of day; during peak periods the lines have more cross-talk (leakage from other conversations) than they have at night.

In addition, the lines of many long-distance services simply won't carry information at 2400 or even 1200 bps. If you have trouble at 2400 bps, try another line, or switch to 1200 or even 300 bps and try again. In some cases, your local phone company can check your line for transmission quality.

Like voice calls, modem-to-modem calls can follow several dialing sequences.

Operator-assisted calls

Whenever possible, you should make modem connections to other computers by purely electronic means: An autodial modem should call an autoanswer modem. In some cases, however, you might need to start the link by calling the operator:

- Dial the phone number manually through an operator.
- When the operator answers, arrange for a collect or credit card call.
- Wait for the connection to be completed.
- Arrange with the other party to switch to data mode.
- Switch together to data mode.

Unfortunately, such switching is difficult or impossible to do with many hardware and software combinations. Both parties need software that can send a command to the modems to switch from voice to data transmission.

Credit card calls

Credit card, or calling card, calls work only with touch-tone dialing, available in most but not all areas:

- Dial 0 and the complete long-distance number.
- After a few rings, you should hear an acknowledgment tone.
- At the tone, dial your credit card number. (If you wait, the special tone stops, and a human operator comes on the line.)

Alternative long-distance services

- Dial the local access number for the service you are using.
- After a few rings, you should hear an acknowledgment tone.
- Dial your account number.
- Dial the complete long-distance number.

To successfully dial a credit card call or alternative long-distance service call with an autodial modem, you'll need to insert commas to pause after dialing the first set of numbers and before any acknowledgment tone begins. If you are dialing from inside a company through a private branch exchange (internal switchboard, or PBX), for example, typing 9,02135551234,,,1222-5551111 will cause the modem to:

- Dial 9 and wait two seconds, allowing time to receive an outside dial tone. (Apple and Hayes modems do not detect dial tones.)
- Dial 0-213-555-1234, indicating to the phone company that you are making a long-distance call with assistance—but will not actually use a human operator.
- Pause six seconds, long enough for the call to clear the phone company's exchange, ring the operator twice, and get an acknowledgment tone.
- Dial 1222-5551111, the credit card number, during the acknowledgment tone. An operator will not come on the line if the card number is dialed at this time.

For the exact sequence of steps, consult your modem and communication software manuals. If you don't have touch-tone service on your local exchange, your modem can still dial the local number with pulses and then, by simple commands, switch to tone dialing to complete the sequence.

Getting It to Work

For many systems, you must send a carriage return or Command-c (Control-c)—sometimes more than one—to elicit a response to your call. (The remote computer can use these characters to determine what speed you are using, among other things.)

Some telephone convenience features, such as call waiting, that produce an audible click can interfere with computer links because a modem will hang up if its carrier signal is interrupted. Even if you have one of the few newer modems that can tolerate and ignore the click, the modem at the other end of your connection may still hang up. If such interference is rare, you might be able to live with it; otherwise you should disconnect these features. Depending on your service, you might be able to disable call waiting temporarily by dialing a special code; after completing your modem call, you can then restore call waiting. Check with your telephone company. A comprehensive communication program should be able to perform these steps automatically.

If you travel with your computer and modem, you might not always be able to connect your modem directly to the phone line. If not, you will have to resort to an acoustic modem with cups, containing a small microphone and speaker, which fits over a standard telephone receiver. Acoustic modems are less reliable because of noise and other problems; always use a direct-connect modem when possible.

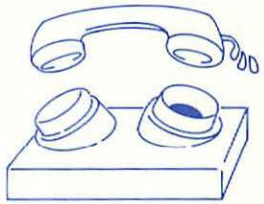

Acoustic cups fit over a standard telephone receiver.

Because communications involve so many steps, finding the exact source of a problem can be difficult. You can test your own hardware and software by calling, say, an electronic mail service, but even a successful connection won't rule out subtle problems in data transfer.

Given how often problems occur in getting modem communications to work, you might find a second phone line handy. With it, you can talk to the person operating the other computer while setting the parameters. Once you have established a link, store all the parameters. The connection should be easy next time, unless line conditions have changed.

COMMUNICATING BETWEEN MACINTOSHES

Moving information between identical computers is usually easy: If the computers are within walking distance, you simply swap disks. Between computers, you can also use a direct wire connection (called a null-modem because no modems are needed), a local area network (see Chapter 20), or telephone lines.

In any case, you should use the same communication software at both ends, set up with identical parameters. On a telephone line with modems, one Mac must be set to originate, the other to answer.

COMPATIBILITY

Communications will not work at all unless the modems at both ends are compatible *and* communication software is compatible with your modem hardware.

Among Modems

Modems follow certain rules—protocols—that specify such details as which audio frequencies will carry the information being sent. All 300-bps modems follow the protocol known as Bell 103, and most modern 1200-bps modems follow the Bell 212A protocol. (Bell Laboratories originally set the standards and gave the protocols their names.) Avoid other 1200-bps protocols, at least in North America.

Newer 2400-bps modems in both Europe and North America use a protocol called V.22 bis. Few manufacturers, however, follow every detail of the V.22 bis recommendations, so both subtle and glaring problems can arise. If you want 2400-bps communication within your company, you should adopt a single modem model.

The 300-bps and 1200-bps modems sold for microcomputers in North America operate at the same signaling frequencies, so they are compatible as long as they are running at the same speed; the most common 2400-bps modems are compatible worldwide.

Nearly all 1200-bps modems operate at 300 bps as well; similarly, many 2400-bps modems can also operate at 1200 and 300 bps. Modems can distinguish among different speeds because each speed is marked by a distinctive carrier signal. Modems transmit this carrier signal continuously while they are running.

Most microcomputer modems operate only asynchronously—that is, a specific signal indicates when each transmitted character starts and stops—although many newer models also support synchronous communication, which is faster because it doesn't need these signals. Synchronous communication is used principally for linking microcomputers to large mainframe computers. (See Chapter 19.)

For successful communications, you must set the correct modem protocol and be sure that compatibility exists at several other levels as well. (See Chapter 19.)

With Software

Most modems contain control circuitry that can dial phone numbers, make a connection, and hang up. To take advantage of these features, your communication software must send suitable instructions to the modem. The most popular instruction set was established by Hayes (Norcross, GA) for its Smartmodem series; Apple and most other modems now sold for microcomputers use this instruction set. Unfortunately, many fail to emulate Hayes modems exactly. The minor differences can cause problems, for some communication programs will run only with a specific modem model. This compatibility problem affects only the commands between the communication software and the modem, not modem-to-modem compatibility.

Well-designed communication software includes installation procedures for different types of modems. Even so, whenever you buy a modem, be sure that it will work with your software; as always, be wary of compatibility claims.

MODEM FEATURES

Beyond the now standard Hayes commands, some modems come with additional features. Callback modems, which do not answer but only return calls to numbers you have stored in their memory, are a good choice for security and privacy, for example; see Chapter 18 for details.

Another feature is a built-in memory buffer for storing incoming messages. The memory is active even when the computer is not turned on. With such a buffer, you don't have to be present when a message comes in; the buffer behaves like a private electronic mailbox. Anchor Automation's Signalman Computer Mailbox (Van Nuys, CA), for example, can store up to 64 KB of incoming information. The Hayes Transet 1000, which is an external memory buffer and processor designed to work with modems, can store up to 256 KB. Such memory buffers make the most sense if you dedicate a phone line to the modem and leave the modem on all the time; in most cases, a commercial electronic mail service is simpler and cheaper.

SOFTWARE FEATURES

All general-purpose communication programs provide for essential operations, such as setting parameters, sending a file from a disk, and storing an incoming file on a disk. The following list outlines potentially valuable features not found in all programs:

- File-transfer protocols. Virtually all communication programs can handle standard text. Most of those for the Macintosh can also transfer programs and graphics by using a special protocol that ensures error-free transmission of any file. The most common of these are Macintosh Xmodem and MacBinary.
- File-transfer time and status. As you begin transmitting a file, many programs tell you how long it will take and show you a progress report.
- Auto-login. When you dial up a remote computer, you usually have to supply a login name and a password to gain access. Some programs let you automate this procedure. Caution: For security reasons you might not want to use an auto-login feature because anyone can read the password off your disk.
- Interactive communication. A recent feature that lets you and another computer user (usually, but not always, on another Macintosh) work simultaneously on a document. The first such Macintosh program is Smartcom II from Hayes, which includes interactive drawing. Two Macintoshes connected together can share the palette and drawing window; anything drawn on one Macintosh also appears on the other. Interactive word processors should also appear in the coming months; with this feature you and an associate across the country will be able to edit a report together.
- Terminal emulation. Many programs can make the Macintosh behave like a DEC (Digital Equipment Corporation) VT52 or VT100 terminal. Some programs also allow Mac to emulate a DATA GENERAL Dasher, Tektronix graphics, or an IBM mainframe terminal. Terminal emulation enables the remote computer to send your Mac short, terminal-specific commands to erase a line, move the cursor around, and so on. In many cases the emulation is not complete, although the missing features might not be critical.
- Session editing. After you are finished communicating, some programs let you go back to edit the entire session and save only important material on disk. The length of a recorded session depends on how much RAM you have available; a few programs record the entire session on disk instead of in RAM.
- Filters. Many computers send a variety of special characters during communications, and in some cases these characters interfere with your

screen display. Inward filters can be set to remove these characters before they reach your screen; outward filters do the reverse, removing characters that might disturb the remote computer. In Macintosh files, the special characters created with the Option key will sometimes cause this kind of interference. Communication programs ought to replace these characters automatically—substituting " for " and ", for example—but no Mac communication program does this yet.

- Background operation under MultiFinder. This means you can start a telecommunication session, begin transferring a long file, and then switch to another application while the communication continues invisibly in the background.

Of the many Mac communication programs, several stand out for their features or ease of use. Among general-purpose programs, we favor MicroPhone, version 2.0 (prerelease, Software Ventures, Berkeley, CA), as one of the best balanced packages. Red Ryder, version 10.3 (Freesoft, St. Louis, MO), lets you set an exceptionally wide range of communication parameters, although no one person is ever likely to need more than a few of them. Smartcom II, version 3.0 (Hayes, Norcross, GA), is unique in permitting interactive drawing.

Two special-purpose communication programs are aimed at specific dial-up services. Desktop Express (Dow Jones Software, Princeton, NJ) greatly simplifies the use of MCI Mail, so you don't have to learn any MCI Mail commands; the program can send and receive mail with no input from you. Although MCI Mail is ordinarily limited to text messages, Desktop Express can use it to exchange binary files with anyone else who is also equipped with Desktop Express or with Lotus Express for the IBM PC (Lotus Development Corporation, Cambridge, MA). CompuServe Navigator (CompuServe, Columbus, OH) is a similar program for CompuServe; it replaces endless menus with an icon-driven interface.

FACSIMILE

Another means of sending information to a remote location via a telephone line is facsimile, or fax. Fax works like a remote copying machine; originals fed into one machine appear, on paper, 20 seconds to a minute later at the remote fax machine. The transmitted information consists of bit-mapped pages, which makes facsimile particularly useful for sending graphics back and forth. Moreover, facsimile transmission is much easier to set up than ordinary modem communications. You merely dial the remote fax machine and send the image; everyone uses the same protocol worldwide.

As of 1988, there were about two and a half million fax machines in use around the world. Most of these were in Japan, where fax has displaced telex and telegrams for sending messages in Chinese characters (kanji); the United States and Canada have about a half million fax machines.

A traditional fax machine is a self-contained unit comprising:

- A scanner to scan paper originals
- A modem with a telephone dialer for transmitting the page image
- A printer to print received images

The resolution and the format of facsimile transmissions follow an international Group 3 standard. (Group 1 and 2 standards pertained to older, slower formats that are now obsolete.) A Group 3 fax image contains 200 dots per inch horizontally and 100 dpi vertically (200 dpi in both directions in "fine" mode). Fax modems run at 9600 bits per second, with automatic reduction to 4800 and 2400 bps if the phone line is noisy or if the receiving fax machine cannot run faster than 4800 bps; they are not compatible with ordinary modems. (For more technical information, see Chapter 19.)

The Mac as Fax

Because fax communicates in a digital format, a microcomputer with the proper accessories can handle some fax functions better than a fax machine. With fax equipment hooked up to a Macintosh, you can exchange two kinds of images with non-Mac fax systems: images initially created on the Mac with a graphics program or images scanned into the Mac and then sent either unchanged or edited with a graphics program. (See Chapter 12.)

For Mac-to-fax communication, you need to add to the Mac an attachment containing a fax-compatible modem and fax software. To get the image from the Mac into the fax machine, the fax software works with a driver in the System file that acts like a printer driver; your Mac application "thinks" it is printing a 200-dpi image. The fax software then takes this 200-dpi image and transmits it through the modem. Images created this way have precisely formed letters and perfect vertical and horizontal lines. They are much cleaner than scanned fax images, which invariably suffer from badly aligned paper, surface faults that create rough edges, and a generally untidy appearance.

If you use a scanner to get an image into your Mac for editing and later transmission by fax, you will run into a problem of mismatched resolutions. Most scanners work at 300 dpi horizontally; to create a fax-compatible 200-dpi image, the fax software ignores every third pixel horizontally. Because scanners physically move a linear scanning array along the page's vertical dimension, however, getting 100 or 200 dpi vertically simply requires the software to increase or decrease the scanning rate.

At the receiving end, an incoming fax image can be displayed on a Mac screen, but because the screen has only 72 dpi, the 200-dpi fax image must be either magnified so that only a small portion of it is visible, or reduced, with considerable loss of detail. A printer makes a much better fax output device than the screen of any microcomputer.

An ImageWriter LQ or a laser printer can print fax images, although mismatched resolution again becomes a problem. An LQ's resolution matches Group 3 reasonably well, but its dot size is a little large. On the 300-dpi LaserWriters, printing can be done with pixel-for-pixel correspondence, which gives you an image two-thirds of original size, or the fax software can alternate single pixels with larger dots made of four pixels, which, despite some jagged edges, minimizes distortion. The ImageWriter and the ImageWriter II do not have sufficient resolution to be satisfactory fax printers.

Turning your Mac into a fax machine carries some operational implications. A typical fax machine is connected to a dedicated telephone line and is left on 24 hours a day. Unless you can afford to devote your Mac solely to fax, this means fax operations must run invisibly in the background so that you can still do normal work in the foreground. Successful background operation will depend on Apple's developing a fully multitasking operating system. (See Chapter 16.) Background operation is unnecessary if you only receive fax images following a voice call to set up the procedure (fax by appointment) or only leave fax running overnight. (With fax attached to a Mac, you can even program the system to make calls automatically at night, when phone charges are low.)

But Does Fax Have a Future?

Although fax technology is digital, it is otherwise unsophisticated. Fax was developed for sending graphics, but most fax images in North America consist only of text. It is far more efficient to send text via electronic mail, which codes characters as characters instead of as bit-mapped pictures of characters. Text messages that arrive via electronic mail can be edited, but fax images can only be edited as bit maps—an extremely cumbersome process. In principle, optical character recognition could convert a fax image into a text file for editing, but OCR of a low-resolution image coming from an unknown source is impractical; be wary of products that claim to do this.

Furthermore, fax images take up considerable storage—typically 30 to 60 KB for a page of text, or 120 KB in fine mode, even with Group 3's data compression scheme. If you devoted a 20 MB hard disk entirely to fax images, you could still store only 170 fine-mode pages. In contrast, a page of text in a computer file that is coded in ASCII or word-processing format takes up only 2 to 4 KB. At 1200 bps, an ASCII page takes less time to transmit via standard modem than a fax page at 9600 bps. If you are sending text to someone with a computer, you are therefore better off sending the text file than a fax image.

For graphics, however, fax makes sense—especially if the recipient doesn't have the same hardware and software you have—because there is no computer equivalent of ASCII for pictures. Even so, if you are sending pictures to another Mac, you should send an object-oriented graphics file that the recipient can edit.

Future fax developments will concentrate on faster transmission and finer resolution, which are easier to get international agreement on than a graphics standard; higher resolution should make OCR practical. But until a universal graphics standard is adopted—if one ever is—fax will remain a reasonable way to send pictures between dissimilar computers, particularly as desktop publishing and its accessories continue to invade today's offices.

TELEX

Telex, another standard means of sending messages via telecommunications, is still widely used around the world, although it is rapidly losing ground to facsimile. It is a primitive service modeled after telegrams; messages can consist of only uppercase characters. Telex communications travel on their own network of wires, separate from the phone line. A Macintosh and other microcomputers can send and receive telex messages through electronic mail services such as MCI Mail.

14: Dealing with Disks

All the information that your Macintosh works with is stored on disks. Information can come on a disk, or you can enter it using the keyboard and mouse, but eventually, everything ends up in files on disks. Thus, the more systematic you are about organizing your disks, the more efficiently you will work. This is true whether you have a hard disk, which can store thousands of files, or only floppies, on which storage space is limited.

Apple's manuals do a good job of explaining the basics of dealing with disks; here we will concentrate on supplementing that information.

Most disk files fall into four broad classes:

- System files: programs and data that are needed when the Mac starts up.
- Utilities: programs and data that help you maintain your system and application disks. They can make operation easier but are not required every time you use a Mac.
- Applications: programs, such as MacPaint or Excel, for doing a particular task.
- Documents: data files that contain the work you create with an application.

If you have a hard disk drive, you can store all these files on a single hard disk. You will need floppy disks only for transporting files to and from other computers and to back up your hard disk. (See Chapter 6.) If, on the other hand, you have nothing but floppy disk drives, you must plan your disk storage carefully to avoid running out of space. (See the box on p. 235, "If You Have Only Floppy Disks.") Unless you plan to use your Mac for only the simplest applications, such as writing nothing but short memos with MacWrite, you should invest in a hard disk drive. Trying to do all your work with floppy disks is too difficult.

THE SYSTEM FOLDER AND SYSTEM FILES

Several files are ordinarily needed to start up a Macintosh:

- The System file: essential to starting up a Mac.
- The Finder: manages files and disks.
- System files (as distinct from the System file), including printing resources, also called printer drivers, which convert information generated by applications into a form that can be sent to a printer; Control Panel files, which contain information about hardware attached to the computer, such as the keyboard, mouse, and video monitor; and INITs, small programs that are executed when the computer starts up.

All these files must be kept in a single folder, usually called the System Folder. In version 6.0 of the Finder, this folder is identified by a small picture of a Mac on it.

The Finder

The Finder creates and organizes the Macintosh Desktop with its disk icons and windows; it also accepts and interprets your commands—to copy files and disks, start applications, open documents, and so on. It is the interface between you and the System file. (See the next section.)

The Finder file need not be on the system disk to start the Mac. If you prefer, you can bypass the Desktop and go directly to a specific application program as soon as you turn on your Mac. To do this, begin at the Desktop and select the application program you want to start in, and then choose Set Startup from the Special menu. Once you have done this, you can remove the Finder file from the System Folder and your system disk, thus saving space on the disk if you are working with floppies. If you quit the application, you will have to restart the computer. With a hard disk, if you start directly in an application but leave the Finder file in the System Folder, you will return to the Desktop when you quit the application.

Several Finder substitutes are available. PowerStation (Software Supply, Sunnyvale, CA) replaces the portion of the Finder that starts up programs; it lets you go quickly from one application or file to another, without threading your way through folders and subfolders. DiskTop, an excellent desk accessory from CE Software (Des Moines, IA), can run many Finder functions, such as searching through folders, copying and renaming files, changing floppy disks, launching new applications, and so forth.

Apple's MultiFinder is an enhancement of the Finder that allows you to load more than one application program into memory. Each application can have one or more windows on the screen, and you can switch among programs by clicking on the various windows. You can also choose applications from the Apple menu or rotate through them by clicking on the program icon at the right edge of the menu bar.

MultiFinder is a powerful addition to the Macintosh, although it had two practical problems when it was introduced. Loading multiple applications requires a lot of RAM, generally more than the one megabyte that was common in 1987. And because

MultiFinder involves many changes to the Macintosh operating system, some Mac applications were not compatible with it. Both problems will be solved in time; as memory prices drop, multimegabyte memory will become common, and updates will fix compatibility problems.

The Finder works closely with the System file and must be compatible with it. Tables 14-1 and 14-2 list Finder/System file configurations for different Mac models; they give Apple's recommendations as of December 1987. The current Finder/System file does not work on early Macintoshes because those machines lack the hardware features and expanded ROMs that are built into newer models. Future Finder/System file updates should work on the Mac Plus, Mac SE, and Mac II, but support for the Plus may stop soon.

TABLE 14-1. SYSTEM SOFTWARE CONFIGURATIONS

	128KB Mac	*Mac XL*	*Mac 512K*	*Mac 512KE*	*Mac Plus*	*Mac SE*	*Mac II*
System 2.0/Finder 4.1	Best choice						
System 3.2/Finder 5.3		Best choice	Best choice	Best choice	OK		
System 4.0/Finder 5.4				OK	OK	OK	
System 4.1/Finder 5.5				OK	OK	OK	OK
System 4.2/Finder 6.0					Best choice	Best choice	Best choice

TABLE 14-2. SYSTEM SOFTWARE CHOICES FOR APPLESHARE FILE SERVER

	Mac 512K	*Mac 512KE*	*Mac Plus*	*Mac SE*	*Mac II*
System 3.3/Finder 5.4	Best choice	Best choice	OK		
System 4.1/Finder 5.5			OK	OK	OK
System 4.2/Finder 6.0			Best choice	Best choice	Best choice

The System File

The System file contains the Macintosh operating system and works in conjunction with the Macintosh ROM; its icon is a Mac labeled System. The System file contains several resources to help you tailor your Mac to your preferences; these include desk accessories, fonts, and FKEYs.

Desk accessories

A desk accessory, or DA program, is usually part of the System file; it is installed or removed with the Font/DA Mover utility program. An installed desk accessory is available any time you choose it from the Apple menu.

Desk accessories do not always have to be installed in the System file; instead, they can be installed in an application program. To do so, simply open the application from within Font/DA Mover and choose Copy. Unlike desk accessories installed in the System file, however, those you install in an application are available only when you are running that application.

With version 4.2 of the System file, you can have a total of 15 desk accessories installed at one time in the System file and an active application combined. This limitation can be circumvented with a utility program called Suitcase (Software Supply, Sunnyvale, CA), which lets you have as many as 500 desk accessories available at once. With Suitcase, you don't need to install any desk accessories in the System file; you can keep them in a specific folder where they are always available under the Apple menu, or you can keep them in other folders where they are available only if you select them with Suitcase. You can change the available DAs at any time, even in the middle of an application.

Desk accessories are supposed to be small programs, but not all of them are. Hundreds of DAs, in dizzying variety, have been written—from games to complete graphics programs. Most are public-domain or shareware programs and are available from user groups. Large desk accessories have been written for nearly every important application, but they will fade in importance as MultiFinder becomes more widespread. With MultiFinder and enough RAM, you won't need desk accessories to do more than one kind of job at a time; you can load several applications into memory and switch among them instantly.

Fonts

Like desk accessories, fonts can be installed with Font/DA Mover either in the System file or in an application program. Four fonts are built into the Mac SE and Mac II ROMs: Chicago 12, Monaco 9, Geneva 9, and Geneva 12 (the numbers indicate point size). With version 4.2 of the System file, you can install a maximum of 256 fonts at a time, although the practical limit is fewer than 200. By definition, each size of each typeface counts as a font (for example, Times 10 and Times 12 are two separate fonts).

The utility Suitcase gives you access to fonts that are stored outside the System file, so a total of about 2000 fonts can be available at a time. Suitcase also lets you group fonts for use with specific tasks, for example, a collection of symbol fonts that you need only when you are running graphics programs. (For more about fonts, see Chapter 12.)

FKEYs

FKEYs (Function KEYs) are short programs installed in the System file that can be called up by pressing Command-Shift-number key (1 to 9 plus 0). Of the ten FKEYs available, Apple predefines four. FKEYs 1 and 2 (Command-Shift-1 and Command-Shift-2) eject floppy disks from the primary and secondary disk drives, respectively; these two programs reside in the Macintosh ROM and are difficult to replace. FKEY 3 produces a MacPaint file of whatever is on the screen at the time, and FKEY 4 prints the active window on an ImageWriter (Command-Shift-Caps Lock-4 prints the entire screen).

FKEYs 5 through 0 are available for user-defined functions, and you can also replace Apple's FKEYs 3 and 4. The easiest way to install or remove an FKEY is through the utility program FKEY Installer (Dreams of the Phoenix, Jacksonville, FL). Many FKEY programs are available, mostly through user groups. They do such varied tasks as printing screen dumps on a LaserWriter, displaying the contents of the Clipboard, or starting a desk accessory.

Which System file are you using?

It's not always easy to keep track of which System file—and therefore which desk accessories, fonts, and FKEYs—you are actually running at a given time. Each time you change application disks, if the disk containing the application also contains a System file, that System file becomes the active one, along with the fonts and desk accessories it contains. The current system disk (sometimes called the current start-up disk) is always displayed as the top disk icon in the upper right corner of the Desktop.

If you are working from a hard disk, simply changing applications—whether to another application on the same hard disk, on a second hard disk, or on a system floppy disk—will not change the System file.

Whether you have a hard disk or floppy disk system, you can always force a change to another System file by pressing the Command and Option keys at the same time and double-clicking on the icon labeled Finder in the System Folder of the disk you want to change to.

On any disk, but particularly on a hard disk, there should be only one System Folder and one System file. If there is more than one of each, you will run into problems ranging from inconsistency to complete system crashes. When you copy files from a floppy disk to a hard disk, take care *never* to copy a System Folder unless you are updating or otherwise maintaining your hard disk.

Printing Resources

Software drivers for a variety of hardware must also be stored in the System Folder. These include printing resources, which control the ImageWriters and the LaserWriters, and the AppleFax modem for sending facsimiles, which is treated as a printer.

The drivers for the ImageWriters and the LaserWriter IISC convert the QuickDraw instructions generated by Macintosh applications into bit maps that the printers can print. The driver for a PostScript LaserWriter, which comes in two parts, loads a setup program into the printer's own computer (LaserPrep) and converts QuickDraw instructions into PostScript (LaserWriter). The driver for the AppleFax modem converts QuickDraw instructions into Group 3 fax format. (See Chapter 19; ordinary modems, which are handled directly by a communication program, do not need a driver.)

A few programs come with their own printing resources. The page-makeup program PageMaker, for example, prepares the LaserWriter with a file called AldusPrep. Consult the manuals for these programs to find out where you need to store their printing resources.

To activate a printer or the AppleFax modem, open the desk accessory Chooser; it displays icons for all the printer drivers in the System Folder. Select the icon for the device you want to use and activate AppleTalk if required, then close the Chooser window. Your system should then be configured for the printer you selected.

INIT Programs

INIT (INITial) programs are those that are executed automatically when the Macintosh starts up. Most are utilities, such as Apple's Easy Access, which makes it easier to use the keyboard with only one hand (great for nursing mothers); public-domain screen savers, which darken the screen if there is no input for a certain fixed length of time; and some electronic mail programs for AppleTalk networks.

CDEV (Control DEVice) programs are a kind of INIT that usually act as hardware drivers or adjust settings for hardware devices. They include Monitors, Key Layout, Keyboard, Mouse, and others, which can be set from the Control Panel desk accessory. When you click on one of the icons in the Control Panel window and change a setting, you are actually modifying a CDEV program.

Not all CDEVs adjust hardware characteristics. Pyro (Software Supply, Sunnyvale, CA), for example, is a screen saver that lets you specify how long to wait before it darkens the screen. QuicKeys (CE Software, Des Moines, IA) is a keyboard enhancer. The primary difference between INITs and CDEVs is that settings defined by an INIT cannot be changed; those defined by CDEVs can.

Both CDEVs and INITs must be in the System Folder before they will work. To activate them, you must first put them in the System Folder, then restart the computer. If you have several INIT and CDEV programs in your System Folder, they will be executed in alphabetical order. To change their order, rename the files; change *apricot* to *zapricot,* for example.

If you try to delete a CDEV or INIT file, you will get a message saying that it is in use. To get around this, move the file out of the System Folder and restart the computer; then delete the file. Although they are stored in the System Folder, CDEVs and INITs are not part of the System file itself.

Preference Files

Many programs can be configured to your preferences. For example, you could set a word-processing program to start every new document in the font and page size you indicate, instead of according to its default specifications. Some applications store your preferences within the program; others create separate files. Some preference files must be in the System Folder; otherwise, they usually must reside in the same folder as the application. Restrictions on where the preference file is stored keep an application from having to search a lot of folders when it starts up.

System Compatibility

From time to time, Apple changes the Macintosh System file and Finder; so far these upgrades have been free to Mac owners. Whenever this happens, though, compatibility among existing software, hardware, and peripherals can be jeopardized; don't be surprised if something doesn't work the first time you install a new System file and Finder. Problems can be particularly severe when the change is major, such as the revisions for the Mac II or MultiFinder. Most successful software publishers also update their software quickly after System and Finder revisions, so compatibility usually suffers for only a short time.

To minimize problems if you are running multiple Macintoshes on a network, use the same version of printing resources for any printer available to the network. It is also a good idea to use the same System and Finder software on all the networked Macs, unless some of them cannot run the most recent System and Finder.

FOLDERS AND FILE STORAGE

Apple has developed two different ways of organizing files on floppy and hard disks. MFS (Macintosh File System) was the earlier scheme and is largely obsolete; the later HFS (Hierarchical File System) is now standard.

On an MFS disk, the disk directory treats all files as equivalent in status, and every file must have a distinct name. Files can be grouped into folders, but this is essentially a cosmetic touch. When you choose Open from the File menu with a word processor, for example, the file dialog box shows every text file, regardless of which folder it is in.

On an HFS disk, files are organized in a branching hierarchy, with folders defining every branch. Two different files can have the same name as long as they are in different folders—that is, on different branches. When you choose Open from the File menu of a word processor, you see only the files in a particular folder. You can move through all the folders, but you see the contents of only one at a time. This structure enables you to group files according to their contents or how you use them.

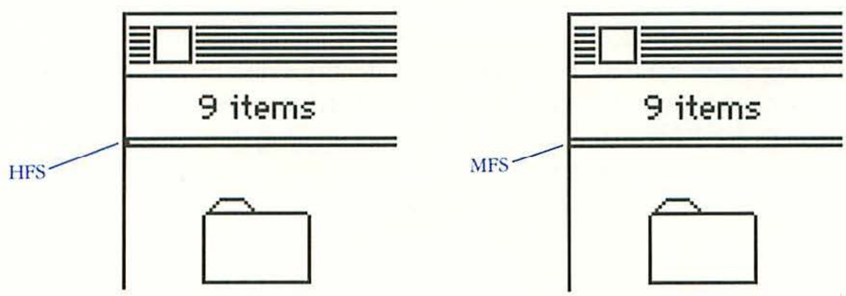

The disk window of an HFS disk has a small dot at the upper left, which an MFS disk window does not have.

MacTree (Software Research Technologies, Laguna Hills, CA) is a utility program that displays the tree structure of an HFS disk; it is particularly useful for keeping track of many files and folders on a hard disk.

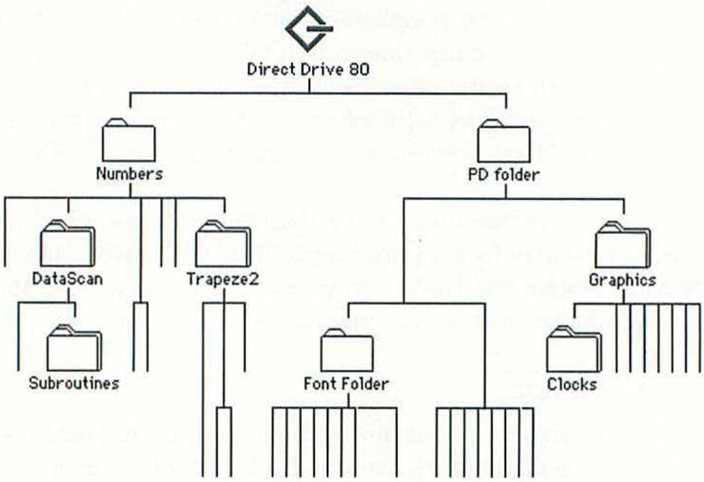

The MacTree utility program displays the branching relationships of the files and folders on an HFS disk.

On all current Macs, hard and double-sided floppy disks are always organized according to HFS. Single-sided floppy disks are arranged according to MFS so that they can be read and written on by 128 KB and 512 KB Macs that might not have the software to deal with HFS disks. In any event, single-sided floppy disks do not usually have the capacity to store enough files for HFS to be necessary. (A single-sided disk can be set up in HFS format if you hold down the Option key while clicking OK in the dialog box to name the disk; or you can hold down the Option key while erasing a disk.)

File Structure

The Mac operating system keeps track of all file and disk names. Usually you select a disk by opening a disk window, but you can also identify a disk by putting a colon between the disk name and the file name. Typing Disk2:Myfile in a Save As dialog box will save Myfile on Disk2. Files on the Macintosh can contain up to three major sections. All files—programs as well as documents—have a header and a resource "fork" or data fork, or both. A fork is merely a portion of the file.

The header

The header, which records basic information such as a file's name, date, and size, is physically stored in the disk-directory tracks. When you put a disk into the Mac, the header is read into the disk-directory buffer in RAM (that's why you can still open a disk window when the disk is ejected). The header specifies whether a file is an application (that is, a program) or a document; it also holds the name of the program that created the file. Thus, if you open a document file, the operating system identifies the program that generated that document and looks for it, first on the same disk and then on other active disks. Application programs are considered their own creators.

The header also identifies the type of data in the file, which is usually unique to a particular application, although some types are not. Many word processors, for example, can create TEXT files (unformatted text), and many graphics programs create PANT (MacPaint) files. (See Table 14-3.) Application programs fall under type APPL.

Other programs besides the creator program can read a document file whose type is not unique. When you try to open a file from within a word processor, for example, the program asks the operating system for all suitable document files. Different word processors might ask for different things: for any text file, for files created only by that word processor, or for files created only on a certain day, although this last is uncommon.

TABLE 14-3. CREATORS AND FILE TYPES

	Creator	*Type*
Finder	MACS	FNDR
System	MACS	ZSYS
Mouse (driver)	cdev	mous
Desk accessories	DMOV	DFIL
Fonts	DMOV	FFIL
DeskTop (a hidden file)	ERIK	FNDR
Microsoft Word	MSWD	WDBN
MacWrite	WORD	MACA
WriteNow	nX^n	nX^d
PageMaker 2.0	ALD2	PUBF
Microsoft Excel	XCEL	XLBN
Filemaker Plus	FMKR	FMKD
Cricket Draw	CRDW	CKDT
Full Paint	PANT	PNTG
Graphic Works	MMCB	CBOK
Adobe Illustrator	ARTY	TEXT
MacDraft	MDFT	DRWG
SuperPaint	SPNT	SPTG
VideoWorks II	MMVW	VW2P

The header does not store the file's icon; icons are stored in the invisible DeskTop file, which assigns them according to a file's type and creator. The DeskTop file is updated along with the header whenever changes are made to a file on a disk.

Finally, the header contains pointers to the physical locations of the resource and data forks, which the system treats as two separate files stored under the same name.

Apple has a utility program called ResEdit (Resource Editor), which can change header information. With it, you can make a program read another program's file, change the icon assignments, and do other housekeeping. ResEdit is widely available from user groups. Caution: ResEdit is for experienced users; mistakes can easily trash your programs and disks.

The forks

The resource fork holds everything that is not conventional data. An application program file usually consists of a header and a resource fork, without a data fork. Resources include the program code, dialog boxes, alert boxes, menus, fonts, icons (for files to be created), and just about anything else that is not a document, although some user data can end up in the resource fork.

The data fork contains mostly document information. A pure text file consists solely of a data fork, but many documents are hybrids because formatting information can go into the resource fork. The distinctions are not absolute; sometimes all the document information is in the resource fork.

KEEPING TRACK OF DISKS

If you use your Macintosh a lot, you will soon find yourself with a fleet of disks that will not keep itself organized. Even finding the latest version of your most needed application can be a challenge. Write the version number and date on all new program disks immediately; otherwise you might later confuse them with updated versions. Programs are available to help you keep track of disks and their contents, such as DiskQuick (Ideaform, Fairfield, IA), which produces indexed and sorted lists of all files on all disks that you process through it.

RAM DISKS

You can set aside some memory so that it behaves as if it were an additional disk drive; doing so creates a RAM "disk," so called because it stores files in random-access memory (RAM). Because it is purely electronic, with none of the moving parts of a real disk drive, a RAM disk is very fast—much faster even than a hard disk. But RAM disks have two important drawbacks: They take away from your memory space, and their contents are ephemeral, disappearing during a power failure or when you turn off the computer. A RAM disk is practical if you have a megabyte or more of RAM; with less, you will be hard put to set some aside.

A RAM disk does not have to use main memory. The RAM allocated to a "disk" can be in an external box connected to the floppy disk drive port or a SCSI port. External RAM disks are expensive and a poor investment, however; it is much better to expand your internal RAM so that it can be put to several uses, RAM disks among them.

DISK CACHES

Like a RAM disk, a disk cache is a portion of RAM that stores disk information. But unlike a RAM disk, a cache does not store specific files; instead, it acts as a buffer between working memory and the disk drives, holding the information most recently read from a disk or written on a disk, thereby speeding up some operations. For example, many programs repeatedly read the same portions of a disk. If that information is already in a disk cache, it is available immediately without the mechanical delays of a disk drive. On the other hand, new information must be read from the physical disk, so there is no speed gain. A well-designed disk cache has a "write through" feature, so any information to be written on the physical disk is written immediately, reducing the risk of losing information if the power fails.

How much speed improvement a disk cache offers depends greatly on the interaction between a specific program and disk files. The optimum size of the cache depends on both the disk-caching software and the application software. Making the cache too small will not produce much speed gain, but in many cases, setting the cache too large will actually slow operation. A few programs incorporate their own disk caches designed to optimize how they use disk files; running these programs with a separate cache will slow everything down. Unfortunately, these programs might not mention anything about disk caches in their documentation; do some testing if you are not sure.

Apple supplies a simple disk cache as part of the Mac System file. You can set the cache size and turn it on or off through the Control Panel; restart the computer to make the settings take effect. In addition, all disk drives—both floppy and hard—have at least a tiny cache of their own; some hard disks have caches as large as 64 KB or more.

A "sticky" disk cache combines the best features of a RAM disk and a disk cache. A sticky cache does everything a disk cache does and also lets you designate files to be loaded into the cache and remain there, regardless of other disk activity, until you turn off the computer.

IF YOU HAVE ONLY FLOPPY DISKS

If your Mac has only floppy disk drives, then you will need to plan carefully to keep from running out of work space on your disks.

Some of your disks—system disks—will contain the system files needed to start your Mac. Others will be document disks that contain only the files you create with your application programs. Depending on space, the application programs themselves can reside on either a system disk or a document disk. If your files are all small enough, you could even put both an application and the documents it generates on a system disk. Indeed, you must do this if you have only one floppy disk drive, but after a little experience you will undoubtedly get another disk drive; when you do, get a hard disk drive, not a second floppy.

To conserve space, put only the files that you need on your working floppy disks. Remove any printing resources that you don't need. If you have an ImageWriter, you will need either the ImageWriter or the AppleTalk ImageWriter driver but not both; remove the LaserWriter and LaserPrep files. If you have a LaserWriter, remove the ImageWriter drivers.

You can usually remove all printing resources from any disk you will not be printing from. When you want to print a file on a disk with no printing resources, save the file, insert a disk that contains the relevant application program and printing resources, and print from that disk.

Some programs, such as early versions of MacPaint, may control the printer directly without using any printing resources; if so, you can remove the ImageWriter file. Others, however, won't run properly without access to the printing resource. Some word processors and page-makeup programs, for example, check the printing resource to set page size.

Fonts and desk accessories take up a significant portion of the System file. Use Font/DA Mover to copy any fonts and unneeded desk accessories from the System file on your working disk(s) to other files or disks; you can always restore any or all of them later. Some fonts are required by the system, and

(continued)

continued

Font/DA Mover will not let you remove them; no desk accessories are required in the System file. The two most important desk accessories are Control Panel and Chooser, and even these can be removed under some circumstances. Control Panel is necessary only for adjusting your system and can be removed once this is done. Chooser is needed to select a printer, but you won't need it on a disk you do not print from.

The Finder file cannot be reduced in size, although you can delete it from your system disk(s) and still start the Mac. Before deleting the Finder from a given system-plus-application disk, select the application and choose Set Startup from the Special menu; when you restart the Mac, you will go directly to that application, bypassing the Finder and Desktop. If you do delete the Finder file, you won't be able to go to the Desktop to copy or delete files and so on unless you restart the computer with a system disk that still has the Finder on it or use a desk accessory such as DiskTop.

You can save a little more space by deleting the help file from programs that let you do so. A help file offers advice about using a program while you are running it, and once you have become familiar with the program, you might not need it. Some programs with a help feature, however, do not have a distinct help file. Caution: Some programs will not start unless their help file is on the disk.

The Clipboard and Note Pad files will use from about a few hundred to a few thousand bytes of space, depending on what you've put in them, and the Scrapbook file potentially more. The Clipboard file is used to store transient Clipboard information if it won't all fit in memory. You can put a Clipboard file in the Trash and erase it by choosing Empty Trash from the Special menu, but the next time you do any work with the Clipboard, the system will create the file again.

The Note Pad file contains entries you have made in the Note Pad desk accessory. This file can also be put in the Trash and erased with Empty Trash.

The Scrapbook file contains all the information you have pasted into the Scrapbook; it can be moved to conserve space. If you move the Scrapbook file to another disk that already has a Scrapbook file, rename the file you are moving (this preserves its contents); you can then delete the file named Scrapbook from the first disk. If you wish to place the contents of the moved file back into the Scrapbook, copy it back to the first disk and rename it Scrapbook.

On a Mac with floppy disks only, you will often need to change system disks when you change applications. If you set up each one with a distinctive Desktop pattern (use the Control Panel desk accessory), you will be able to distinguish them quickly.

15: Programming Languages

Programming languages inspire heated discussions. Each language has its strengths and weaknesses; each also has strong advocates who insist that no other language is worthy of attention. Recognizing that others may not agree with the opinions that follow, we offer this brief guide to computer languages in general.

DO YOU NEED TO PROGRAM AT ALL?

For most people, the answer is no. Increasingly sophisticated application programs will fill all common requirements, including most jobs previously accomplished by programming. So the majority of microcomputer users will never write programs—at least not in the traditional sense. When you work with a spreadsheet or database, you are doing a form of programming, but you are using a very high-level language designed for a specific task. Traditional programming would use a general-purpose language to build up the steps required to perform the task.

Certainly no microcomputer user should begin with programming; everyone should learn applications first. But if you need functions that packaged software cannot provide, or if you are simply curious, by all means learn to write your own programs.

WHAT LANGUAGE SHOULD YOU USE?

For hobbyists, the choice is between using HyperTalk, the programming language built into Apple's HyperCard database program (see Chapter 11), or using a more conventional computer language. HyperTalk lets you build Mac-style programs quickly, but mastering a conventional language will ultimately give you more flexibility and power. HyperTalk aside, the choice of language is not usually critical. The best choices are the most popular languages, such as BASIC and Pascal, since you can get help from many books, magazines, and friends.

Children may find programming interesting and absorbing. Parents should not, however, assume that skills in today's computer languages will ensure a child's employability in coming years. By the time today's children become adults, vastly more powerful computers will have changed programming greatly. There is even the common, although dubious, argument that learning traditional languages will inhibit fluent use of more powerful programming tools.

For children, the best language to begin with is probably whatever language is taught at school—BASIC and Logo are the most common. For high-school students studying computing, the College Entrance Examination Board chose Pascal as the standard language for its Advanced Placement examinations. In the future, the Board may also accept other languages.

In companies, the decision of whether to program and which language to use can be critical. Poor choices, even if they yield quick results, can and do lead to years of frustration later. In many corporations, traditional computer languages are already a severe barrier to progress. Programs written in those languages are increasingly difficult to maintain and improve. Data-processing departments have begun spending more time trying to maintain existing programs than writing new ones. Getting out of this fix requires abandoning traditional languages and finding new ways of writing programs. For an excellent discussion of the current situation in corporate data-processing departments, see *An Information Systems Manifesto* by James Martin (Prentice-Hall, 1984).

AN OVERVIEW OF LANGUAGES

Computer languages fall—more or less—into two general categories: high level and low level. Although the boundary is not well defined, high-level languages, such as BASIC, Pascal, and Logo, contain elements of English in their instructions (source code), which makes them easier for people to use. Low-level, or machine, language, in contrast, deals directly with the microprocessor and consists almost entirely of cryptic abbreviations or numbers. The distinctions among languages are rarely absolute, though, because the successful ones evolve over time, gaining attributes of other languages.

Languages can also be either imperative or object oriented. The most common high-level languages are imperative, consisting mostly of commands that manipulate defined objects in specific ways. For example, you might begin with the declaration:

```
A, B, and C are numbers.
```

and then give the instruction

```
C = A + B
```

and the language will manipulate the numbers accordingly.

In contrast, object-oriented languages send messages to objects; the response depends on the object. To add two numbers, your instruction follows this form:

Message to object A: add yourself to object B.

When object A receives the "add yourself to" message, it is responsible for performing the instruction or complaining if it cannot; the programmer is freed from the detail work that the objects carry out. Mixing two different types of object can create new objects that possess the attributes of both types, so object-oriented languages are suited to defining new procedures and recursive (repetitive) functions in which each step depends on the result of the previous one.

Language Structure

The computer instructions, or source code, written in early languages all come out in a flat, linear sequence—a long list of statements executed in order. All lines are or can be numbered, and the order of execution can be changed by instructions to go to another line. Changes in one part of the program often force changes in many other parts. The complex interrelationships of instructions within these programs inspired the term *spaghetti code* because figuring out how the programs work is so difficult.

Most modern languages, however, use a structured form. Programs written in these languages are divided into blocks, or logical subunits. Each block is a complete procedure and can be changed without affecting any other procedure. Blocks are usually set up by indentations in the source code, somewhat like an outline. The layout is supposed to encourage the programmer to think in an orderly way.

INTERPRETERS AND COMPILERS

Interpreters and compilers take high-level language instructions and turn them into low-level machine instructions. A language interpreter takes one line of high-level instructions and converts it to the low-level form the microprocessor can understand. The microprocessor carries out that line of instructions and hands control back to the interpreter, which converts the next line, and so forth.

Because it works a line at a time, an interpreter is less efficient than a compiler, which translates the entire high-level language program into a machine-language program before carrying it out. A compiled program will usually run much faster than an interpreted one. But developing and debugging a program is much faster with an interpreter because you don't have to stop and recompile the whole program to find out if it works.

Many languages operate in a fashion that is somewhere between an interpreter and a compiler. Some languages permit a program to contain subroutines (subsections of the program) that have been compiled from another language.

For a useful introduction to computer languages, see *Computer Languages: A Guide for the Perplexed* by Naomi S. Baron (Anchor/Doubleday, 1986). The *McGraw-Hill*

Encyclopedia of Electronics and Computers contains a concise but highly technical essay on computer languages.

LANGUAGES FOR THE MACINTOSH

A steady stream of languages has appeared for the Macintosh. The languages described briefly here do not exhaust the list of possibilities. (See the end of this chapter for a comparison of source codes in several different languages.)

BASIC

BASIC (Beginners' All-purpose Symbolic Instruction Code) is the most common microcomputer language and remains a good choice for short, relatively straightforward jobs. BASIC was originally written in 1964 by John Kemeny and Thomas Kurtz at Dartmouth College as a simplified form of FORTRAN. BASIC is one of the easier languages to learn—easier, for example, than either Pascal or Logo.

In its original, unstructured form, BASIC works somewhat like stream of consciousness: It is quick and intuitive but often a little untidy in actual execution. You will probably find it difficult to follow a BASIC program written by someone else, or even a program you wrote some time ago, unless unusually clear annotations are included in the source code.

BASIC has been a moderately well-specified language; industry standards exist for the instructions and syntax, but now there are two virtually distinct forms: the old BASIC and the new, structured versions.

Microsoft QuickBASIC and True BASIC are similar versions, or dialects, for the Macintosh. With either of them you can write a program in one window and see its output in another. However, these and other BASICs are not compatible; you cannot write a program in one dialect and run it in another. Both are much easier to use than earlier BASICs. Unlike earlier BASICs, QuickBASIC and True BASIC use a structured design; they have no line numbering, and indentations mark the beginning and end of procedures in the source code, much like Pascal. These changes answer the most severe criticisms leveled against BASIC in the past. QuickBASIC includes both an interpreter and a compiler.

With the addition of structured programming, QuickBASIC has changed greatly from earlier Microsoft BASICs developed for the IBM PC and the Apple II. It can, however, run programs—including those written in AppleSoft BASIC, which is built into the Apple II—written for the earlier versions. It remains compatible with MBASIC for other computers except for instructions that are specific to a particular microcomputer design—mainly variations in graphics and sound generation. You can thus run or adapt many programs from earlier books and magazine articles and also write new programs in the structured form. The structured version is appearing for other microcomputers too. If you look at books on BASIC, check to see which version is described.

True BASIC was also written by Kemeny and Kurtz. It incorporates their current thoughts about the language and represents their first effort for microcomputers. A dialect of True BASIC that is compatible with Macintosh source code is available for the IBM PC; even graphics programs can run without change. True BASIC compiles to intermediate code; a run-time module is available.

Pascal

Pascal (named after seventeenth-century mathematician Blaise Pascal) was designed by Niklaus Wirth, who based it in part on Algol, a language more popular in Europe than in North America. Pascal comes much closer than earlier languages to the academic concept of how a good computer language should work. Pascal popularized structured programming; it has been a better choice than BASIC for long programs, but its advantage is fading as BASIC becomes more structured.

Pascal standards are spelled out by the International Standards Organization (ISO). Several dialects exist, but the presence of clear standards keeps them from straying too far. With Pascal programming experience, you have a head start in learning C, now the most popular developmental language for commercial microcomputer programs.

Several Pascals are available for the Mac. TurboPascal (Borland International, Scotts Valley, CA) and Lightspeed Pascal (Think Technologies, Bedford, MA) appeal to the widest range of programmers. MPW Pascal, developed by Apple and distributed by Apple Programmer's and Developer's Association (Renton, WA) is part of a complete application development package.

Modula-2

Niklaus Wirth wrote Pascal as a teaching language; he went on to develop Modula-2 (*Modu*lar *la*nguage-2), which is similar, as his "real" language. Modula-2 cleans up rough edges in Pascal and adds missing elements, such as the ability to compile a program in separate modules. Modula-2 also offers features interesting to advanced programmers. For example, it supports a form of concurrent programming, in which the execution of one segment of a program can be halted to allow another segment to run, so it can serve as a real-time control language. With real-time control, a program can be interrupted to respond immediately to specific events, such as a message from a peripheral device. Modula-2 uses a compiler to create an intermediate code that is then interpreted to run the program.

C

C (written after the language B was written) is a middle-level language developed by Dennis Ritchie at Bell Telephone Laboratories for writing the UNIX operating system. C is structured and uses memory efficiently, and finished programs run quickly, but it is fairly difficult to learn. It is currently the most popular language for serious microcomputer software development. On the Macintosh, you can choose from many different C compilers. Fortunately, comparisons are beyond the scope of this book.

Logo

Logo (from the Greek *logos,* meaning word or thought) was originally written by Don Bobrow and a group at MIT headed by Seymour Papert. It is best known as a language for children because its graphics are very easy to learn, but it is a full programming language. For complex programs, Logo is harder to use than BASIC or Pascal; progressing beyond the simplest steps is difficult. Because Logo is an extendable language (you can define new procedures within the language), there is no standard form.

LISP

Experience with Logo can be applied to LISP (LISt Processing), a language developed specifically to deal with lists and recursive definitions (nearly circular definitions that build upon themselves), which are hard to handle with traditional languages. You can define and manipulate lists of numbers or text or other lists. With long text strings to compare and recombine, LISP uses up memory quickly; several megabytes of RAM are not too much to have. LISP was written for, and remains largely limited to, academic and industrial research into artificial intelligence. LISP was the precursor language of Logo.

Forth

Forth was originally written by Chuck Moore to control radio telescope-positioning motors. It was initially called Fourth, for fourth-generation language, but the IBM 1130 computer used for its development would accept only five-character names.

Forth works with a stack of variables organized in reverse-Polish notation, a nonalgebraic form of mathematical expression in which the operator (plus, minus, and so on) follows the numbers it affects: For example, 23 + 2 is written as 23 2 +. (Reverse-Polish notation is the same notation used on a Hewlett-Packard calculator.) Programming essentially consists of keeping track of where you are in the stack; since the stack can be many thousands of variables deep, this can get complicated. For the skilled programmer, however, Forth program development is quick, especially for procedures involving many numbers. Forth programs are hard to read and modify but use memory efficiently. Forth is moderately well defined.

COBOL

Since 1960, COBOL (COmmon Business Oriented Language) has been the most widespread business language on large computers. Unlike most other early languages, which were designed to deal with numbers, COBOL was designed to deal with text and records; it looks more like English than other programming languages do. Its critics are legion, charging that COBOL programs are hard to maintain and modify and that use of the language continues out of inertia. Consider COBOL only if you work with existing programs that are not available in another language.

FORTRAN

A language from the 1950s, FORTRAN (FORmula TRANslation), like COBOL, has also seen its time come and go. It has been in a slow, steady decline for more than a decade. FORTRAN was originally developed for solving mathematical problems, and it remains entrenched in universities and research centers as the main language for scientific computation. As with COBOL, use FORTRAN only if you must work with existing software.

APL

Kenneth Iverson, working first at Harvard and then at IBM, developed APL (A Programming Language) in 1960 for manipulating arrays of numbers. Especially good for numbers arranged in matrices, it is used for engineering and scientific tasks. APL source code is unusually compact; the only control statement is a GOTO command. The source code contains odd-looking symbols, which make APL programs distinctive and also inspire the jibe that APL is a write-only language: Once you have written a program, no one can figure out how it works. Long an interpretive language for mainframe computers, APL has appeared on microcomputers in recent years.

Graphical Languages

A new class of programming languages may emerge on the Macintosh: languages that use icons and flow diagrams instead of lines of source code. The first such language to appear is VIP, or Visual Interactive Programming (Mainstay, Agoura Hills, CA). For years, instructors have taught programming students to begin with a flow diagram before writing a program. Hardly anyone ever does, but with VIP, the flow diagram itself is the program.

HyperTalk

The structured language built into Apple's HyperCard, HyperTalk contains commands to manipulate HyperCard's stacks and graphical interface. Although it is easier to use than most conventional languages, HyperTalk does not make programming accessible to the rest of us; it requires the same careful planning and discipline as any other computer language.

Assembly Language

Because high-level languages are so different from the machine code that a central processing unit (CPU) can execute, programs written in high-level languages often can't take full advantage of a CPU's power, and they usually run slowly. Programs written directly in machine language would optimize speed and efficiency, but machine code consists purely of numbers and is much too tedious and difficult for most people to write. Assembly language, a low-level language close to machine language but easier for humans to work with, offers one solution by replacing the numeric codes with mnemonics. Even so, only the most determined programmers write in assembly language.

Disassemblers take existing program code and display the assembly-language equivalent. MacNosy from Jasik Designs (Menlo Park, CA) is an interesting disassembler that automatically follows logical paths. It works on ROM programs as well as on conventional disk-based programs; it's for experienced assembly-language programmers only.

For adequate speed and flexibility, most commercial programs for the Macintosh are written in a combination of compiled code (most often C) and assembly language. Apple has not yet specified a standard format for compiled code, so you might not be able to link modules written in different languages. However, companies that offer several languages often ensure that modules written in any of their languages will be compatible.

THE MACINTOSH TOOLBOX

For effective operation on a Mac, any program should be able to use the Mac Toolbox, which consists of the programs in ROM. But programming-language manuals contain, at best, only sketchy information about the Toolbox. The standard reference comes from Apple as *Inside Macintosh* (Addison-Wesley, 1985, with later supplements). For more accessible discussions, see the two volumes of *Macintosh Revealed,* second edition, by Stephen Chernicoff (Sams/Hayden, 1987).

DEMONSTRATION PROGRAMS

The following simple programs in a variety of languages do exactly the same thing: They read a number, A, from a disk file, and then they read a number, B, from the keyboard. Each program adds the two numbers to find C and then draws a horizontal line the length of C.

The authors of these programs were asked to make them comparable, rather than fancy. These samples hardly show the nuances of the languages used, but you can compare them and get a glimpse of them in action.

```
' Drawline--Microsoft Basic 2.0 program by Philip W. Marshall
OPEN "Drawline Data File" FOR INPUT AS #1
INPUT #1, A
INPUT "Enter a number", B
C = A + B
LINE (0,0) - (C,0)
END
```

```
! True BASIC program by Brig Elliot
!
open #1:name "MyFile"
input #1:a
input b
let c = a+b
plot 0,1;c,1
end
```

```
program Drawline;
{ Macintosh Pascal program by Philip W. Marshall }
 var
 A, B, C : integer;
 DiskFile : file of integer;

begin
 reset(DiskFIle, 'Drawline File');
 read(DiskFIle, A);
 write('Enter a number: ');
 readln(B);
 C := A + B;
 moveto(0, 0);
 lineto(C, 0);
end.
```

```
MODULE Drawline;
(* Modula-2 program by Donald L. Cohn *)
FROM InOut IMPORT ReadInt, ClearScreen, WriteString, WriteLn,
      OpenInput, CloseInput;
FROM QuickDraw1 IMPORT PenSize, MoveTo, LineTo;

VAR
 a, b, c : INTEGER;

BEGIN
 ClearScreen;
 WriteString("Respond to the prompt below with the name of a file");
 WriteLn;
 WriteString("containing an integer.");
 WriteLn;
 OpenInput("TEXT");
 ReadInt(a);
 CloseInput;
 WriteString("Give an integer. ");
 ReadInt(b);
 c := a + b;
 ClearScreen;
 MoveTo(0, 0);
 PenSize(1, 4);
 LineTo(c, 0)
END Drawline.
```

```
10 REM Microsoft Basic 1.0 program by Philip W. Marshall
20 OPEN "Drawline Data File" FOR INPUT AS #1
30 INPUT #1,A
40 INPUT "Enter a Number: ",B
50 LET C = A + B
60 LINE (0,0) - (C,0)
70 END
```

```
/* C program by Greg O'Brien */
#define SMALL_MEM
#include     <inits.h>
#include     <quickdraw.h>
#include     <dialog.h>
#include     <pb.h>
#define NIL (char *)0

char  numbuf[16] ;
char  volName[16] ;
long  count ;
int         a,b,c ;
int         item;
short type = 0 ;
short vRefNum ,
            fd ;
DialogPtr dp;
Rect  box ;
Handle      itemPtr = 0 ;

main()
{

      InitGraf(&thePort);
      InitWindows();
      TEInit() ;
      InitDialogs((ProcPtr)NIL);
      SetCursor(&arrow) ;

      /* Read a number from a file */
      GetVol(volName,&vRefNum) ;
      FSOpen("\Pdatafile",vRefNum,&fd) ;
      count = 3 ;
      FSRead(fd,&count,numbuf) ;
      FSClose(fd) ;
      a = atoi(numbuf) ;

      /* Read a number from the keyboard */
      dp = GetNewDialog(7, (Ptr)0, (WindowPtr)-1);
      /* Read in value from text box */
      do {
            ModalDialog((ProcPtr)NIL, &item);
      } while (item != 1);
      GetDItem(dp,3,&type,&itemPtr,&box) ;
      GetIText(itemPtr,numbuf) ;
      DisposDialog(dp);
      b = atoi(&numbuf[1]) ;
      c = a + b ;

      /* Draw a line of pixels on the screen */
      MoveTo(0,0) ;
      LineTo(c,0) ;

}
```

(continued)

continued

```
TO DRAWLINE
; Mac Logo program by Kerry E. Lynn
; Draws a horizontal line of specified length in upper
; left hand corner of graphics window.
Clean
SetWSize "Graphics [200 200]
Open "MyFile "#Disk
SetRead "MyFile
Make "A ReadWord
Close "MyFile
Type [Enter a positive integer :]
Make "B ReadWord
Make "C :A + :B
PenUp SetPos [-100 100] PenDown
SetHeading 90
Forward :C
END
```

```
(defun drawline (&aux a)
  ;;ExperLISP program by Dexter Pratt
  (moveto -230 -75) ;top left of default graphics window
  (with_open_read "ExperDemo:disk-data"
        (setq a (read)))
  (print "input a value for B -> ")
  (line (+ (read) a) 0)
```

```
SCR # 15          "MyBlocks"   04/16/85   03:00:16 PM
( MacFORTH program by Kerry E. Lynn )          ( 041685 KEL)
VARIABLE MyFileNum
VARIABLE A        VARIABLE B        VARIABLE C
: ReadFile ( --- | read an int from file "MyFile" into var "A" )
  NEXT.FCB MyFileNum !       ( next available file number )
  " MyFile" MyFileNum @ ASSIGN  ?FILE.ERROR
  MyFileNum @ OPEN             ?FILE.ERROR
  4 MyFileNum @ SET.REC.LEN  ( MacFORTH ints are 4 bytes long )
  A 0 MyFileNum @ READ.FIXED  ( read record 0 into var "A" )
  MyFileNum @ DUP CLOSE REMOVE ;
-->

SCR # 16          "MyBlocks"   04/16/85   03:00:28 PM
( MacFORTH program by Kerry E. Lynn )          ( 041685 KEL)
: ReadKB ( -- n | return an int from keybd, no range checking )
  0 >IN ! 0 BLK ! QUERY            ( get string from keybd )
  BL WORD BL OVER COUNT + C! NUMBER ; ( convert it to a number )
: Drawline ( --- | draw a line in corner of interpreter window )
  GINIT   PAGE
  ReadFile
  ." Enter a positive integer : "   ReadKB B !
  A @ B @ + C !
  0 0 MOVE.TO  C @ 0 DRAW.TO ;
```

```
Add section with COBOL code:
Identification Division.
Program-ID. Drawline.
* COBOL program by Micro Focus Inc.
Environment Division.
Input-Output Section.
File-Control.
    Select MyFile Assign "Example"
    Organization Sequential.
Data Division.
File Section.
FD MyFile.
01 A Pic 9999.
Working-Storage Section
77 B Pic 9999.
77 C Pic 9999 Comp.
77 Xpos  Pic 9(4) Comp Value 0.
77 Ypos  Pic 9(4) Comp Value 0.
77 ZerO  Pic 9(4) Comp Value 0.
77 MoveTo Pic X(4) Value X"A893000A".
77 Liine Pic X(4) Value X"A892000A".
78 MacRom Value X"A5".
Procedure Division
    Open Input MyFile.
    Read MyFile.
    Accept B.
    Add A, B Giving C.
    Call MacRom Using MoveTo, Xpos, Ypos.
    Call MacRom Using Liine, C, ZerO
    Stop Run
```

```
      Program Drawline
!     Fortran 77 program (MacFORTRAN) by Paul S. Linsay
!
      Include Toolbx.par
      Integer*4 A,B,C
!
      Open(UNIT=1,FILE='Drawline.dat',STATUS='OLD')
      Read (1,'(i2)') A
      Type 'Enter B '
      Accept B
      C=A+B
      Call Toolbx(MOVETO,5,25)
      Call Toolbx(LINE,C,0)
      Pause
      Close(1)
      End
```

```
    ∇DRAWLIN
     ⍝ APL PROGRAM BY BILL WESTLAND
     'DRAWFILE' ⎕FHTIE 1 ⋄ A←⍞ ⎕FREAD 1 ⋄ ⎕FUNTIE 1
     B←⎕
     C←A+B
     MOVETO 1,1
     LINETO 1,C
    ∇
```

Now, here's a segment of an assembly-language program that performs the same simple task (the first 34 lines have been omitted):

```
; 68000 assembler program by Greg O'Brien
;     Read a number from a file
;
      pea   vRefNum(A5)      ; where to store the volume ref. number
      pea   volName(A5)      ; address of volume name string
      _GetVol                ; get volume reference number
      add.w #8,a7            ; clear arguments from the stack

      pea   fd(A5)           ; where to store file descriptor
      move.w   vRefNum(A5),-(a7); which disk volume to read
      pea   Filename(A5)     ; name of file to open
      jsr   _FSOpen          ; open the file
      add.w #10,a7           ; clear arguments from the stack

      move.l   #3,count(A5)     ; number of characters to read
      pea   numbuf(A5); where to place characters that are read
      pea   count(A5)   ; how many characters to read
      move.w   fd(A5),-(a7); specify file to read from
      jsr   _FSRead          ; read from file
      add.w #10,a7           ; clear arguments from the stack

      move.w   fd(A5),-(a7); specify file to close
      jsr   _FSClose    ; close the file
      add.w #2,a7       ; clear arguments from the stack

      pea   numbuf(A5); where to get string
      jsr   atoi        ; convert ascii string to binary integer
      add.w #4,a7       ; clear arguments from the stack
      move.w   d0,A(A5)    ; store value in A
;
```

By now, you get the idea. The actual program continues for another 42 lines.

And finally, here's the machine-language code. Again, length decrees that lines be omitted; this time, 437 of them:

```
0100111011010101
0000011001100110
0011011100110111
0010111001110010
0110110000100000
0110111001101111
0111010000100000
0110011001101111
0111010101101110
0110010000001101
```

MICROSOFT BASIC NOTES FOR PROGRAMMERS

If you want to move a program written in AppleSoft BASIC to Microsoft QuickBASIC on the Macintosh, you must write replacement graphics instructions for the following Apple II–specific statements: HTAB, VTAB, HGR, GR, HLIN, VLIN, HPLOT, HCOLOR, COLOR, HGR2, and PDL. Any PEEKs and POKEs in AppleSoft BASIC will not work on the Macintosh.

If you want to move a BASIC program from the IBM PC to the Mac, these statements in IBM BASICA (version 3.2) are not supported by Macintosh QuickBASIC: BLOAD, BSAVE, COLOR, COM(n) ON/OFF/STOP, DEF SEG, DEF USR, DRAW, KEY, LOCATE, MOTOR, ON (event) GOSUB, OUT, PAINT, PEN, PLAY, SCREEN, SOUND, STRIG, and WAIT. Similarly, the following IBM BASICA intrinsic functions are not present in QuickBASIC for the Mac: CRSLIN, INP, SCREEN, STICK, and USR. And finally, these machine-specific statements or intrinsics will need modification: CALL, CLEAR, PEEK, POKE, and VARPTR.

16: More About Software

Behind every piece of software is a software developer. And behind every version that actually reaches the marketplace are several other versions in various stages of disarray. Alpha software is equivalent to a first draft; it is distributed to the programmer's coworkers so that they can find bugs. Beta software is like a further draft, considered good enough for outside review but not yet ready for publication; it is distributed to nonpaying customers (like journalists) so that they can find bugs. Released versions are equivalent to published books in the bookstore; they are distributed to paying customers, so that they can also find bugs...

This chapter contains a grab bag of equally vital information about software, including discussions of user interfaces, choosing software, moving information between Macintosh programs, copy protection, and operating systems.

THE USER INTERFACE

While using a microcomputer, you are constantly making choices—to open a file, close a file, insert or delete text, and so on. Software design must therefore address two key questions:

- How does the computer show you what choices are available?
- How do you tell the computer which choice you have made?

The solutions to these problems make up the user interface.

Conventional Interfaces

All conventional interfaces are based solely on text because many computers cannot handle graphics well. Although most Macintosh programs do not use these interfaces, you will encounter them if you use other computers, if you use your Macintosh as a terminal for other computers, or if you run software originally designed for other microcomputers on your Macintosh.

Command-line interfaces rely on commands typed in by the user. The screen shows only a prompt, such as A> or $, indicating that the computer is ready to accept a command. There is no other information; you must know what to type next, and you must type it exactly, so these interfaces are the hardest to learn. On the other hand, they are the easiest to program and offer great flexibility because they can be set up to recognize any keyboard entry as a command. Operating systems that use command-line interfaces include CP/M-80 and MS-DOS.

Menu-initial interfaces present you with a series of single-letter abbreviations that stand for possible choices. In a display such as B C D E F G I M, you must know which letter represents the choice you want. Often, choosing one letter from the menu produces a second-level menu with more letters prompting further choices. Still fairly cryptic, menu initials are nevertheless a little easier to use than command lines. Programs using menu-initial interfaces include the early versions of VisiCalc and SuperCalc.

Menu-word interfaces give you a list of words instead of initials to choose from. Again, choosing an item may take you to a second-level menu with more words indicating more choices. Although less cryptic than menu initials, menu words take up more space on the screen. Menu-word interfaces are sometimes called moving bar interfaces because the selected word is highlighted, making it look like a bar. Programs using menu-word interfaces include non-Macintosh Multiplan, Lotus 1-2-3, and SYMPHONY.

Menu-driven interfaces present you with a complete menu from which to choose. The menu takes up the full screen; with each choice, a new menu appears, until you have made all the necessary choices. Menu-driven interfaces are good for the novice because they can display full instructions, but they are exasperating for experienced users because it takes so much time to wade through all the menus. Examples of menu-driven interfaces include Wang word processors and many traditional accounting packages.

A New Solution: The Visual Interface

Beginning in the 1960s, computer research groups began looking for ways around the user-interface logjam, setting as key requirements ease of use, consistency, and familiarity. Xerox Palo Alto Research Center (PARC) developed an interface incorporating a mouse, icons, and pull-down menus, which has been the most successful one yet. This interface can offer many choices without interfering with work in progress. You choose a menu category from a menu bar and "pull down" the menu with a mouse. The work area is obscured only while you are choosing from the menu. (The menu bar could have been designed to run along the screen bottom and pop up, but the menus would then cover up the most likely active area—the bottom of your document.)

The PARC interface has had tremendous success, and all new microcomputers are using a form of it. Apple built the PARC interface into its Lisa and Macintosh. Companies as diverse as AT&T, Atari, and Commodore have adopted it. Software such as Digital Research's GEM, Microsoft Windows, and Presentation Manager for OS/2 install a PARC interface on IBM PCs. The most recent versions (3.4 and later) of MS-DOS include an optional "shell" that has some features of a PARC interface.

The widespread incorporation of this visual interface means that the microcomputer industry is adopting a standard by consensus. When you learn to use the interface on one machine, you have pretty much learned how to use it on another. You might not know all the details of operating different machines, but major procedures such as starting a program, editing text, and saving a file will be nearly the same. Programs with other interfaces will become increasingly harder to sell.

Nevertheless, many programs designed in an earlier era will linger as the industry changes. A few companies have adapted their programs to run on the Macintosh exactly as they do on other microcomputers. There is little point in using these programs unless they meet a specific need; even then, you should probably replace the program when a Mac-style equivalent appears.

Other software companies will rewrite their programs to take advantage of the visual interface. You will see the Macintosh interface, but the programs will still retain their unique features and capabilities. In many cases, the Mac version will have the same file format as versions for other computers, enabling the Mac to exchange information with the other computers via modem; in others there will be a file-conversion program. (See Chapter 21 for more about exchanging information between computers.)

Software for the Mac doesn't have to incorporate the Mac interface, however; a programmer can write an entirely different interface or, more likely, add new elements to the present one. Such changes can be confusing if you're used to the standard Mac interface, so developers ought to avoid them without good reason.

Is Mac's visual interface the optimum one? No one really knows. But the development of interfaces should be far greater during the next ten years than during the past thirty, so perhaps we will find out.

CHOOSING SOFTWARE

For important computer applications, many software packages compete for your attention, and most of them include essential features. So how do you choose among them?

The choice may depend less on features than on how well the program meshes with your other work. Selecting among competing packages rests largely on such considerations as personal preference, analysis of your needs, and, increasingly, compatibility among programs.

But even before considering which software you need, you must ask yourself another question.

Do You Need the Software?

Longtime computer users can show you shelves full of programs they bought but now ignore. Lots of ideas can sound terrific at first and then turn out to be impractical.

The classic example is a calendar program. Many calendar programs let you enter appointments and notes and have the computer beep at you on cue. A new crop of calendar programs is based on HyperCard. But none of them works as well as a small pocket calendar and the alarm on your wristwatch, which are available at all times, not only when you are at your computer. How do you enter a last-minute change in an appointment when you're in a phone booth? Some people—doctors, for example—do need comprehensive scheduling software, but the common calendar programs aren't good enough for such needs.

Criteria to Think About

If your main goal is to get work done rather than play with a computer, try to choose software that minimizes the length of time you spend learning to use it. Programs styled after another package you have already mastered should be your first choice. Although most Macintosh programs use a common interface, details such as command-key assignments and file compatibility vary considerably.

Consider the following as well:

- Growth. Will the software in question grow with your needs? For example, can you increase the complexity of your models in a spreadsheet program? Can you incorporate into your spreadsheet information produced by other users? If a program cannot expand with your changing needs, it could be a dead end, forcing you to start over with a new program. A well-designed, complex program will let you use its simpler features without fuss; you can start using the more sophisticated features when you need them.
- Manuals. Is the program's manual effectively organized and well written? Does it have a decent index? Can you find information quickly? Does it use the same wording for operations as Apple's manuals? Apple provides authors with a writer's style guide, so anyone who writes about Macintosh software should be using the same terms with the same meanings—*choosing* a menu item versus *selecting* some text, for example.
- Support. Can the dealer from whom you bought the software help you adequately if you have a problem? How about a friend or business associate? Does it look as if the software company will remain in business? For game programs, it may not matter; for a major database program, it is vital.
- Compatibility. Will the program work with your other software? Can it transfer files to and from your other programs and other computers? With many programs, you need to massage files quite a bit before the information can be used in another program.

MOVING INFORMATION BETWEEN MAC PROGRAMS

Macintosh programs can exchange information in five ways:

- Through the Clipboard and Scrapbook
- As "live," or linked, information between program components in an integrated package
- With a screen dump
- Through document files that more than one program can read
- Through special links between two independent programs

The Clipboard and Scrapbook

Whenever you put something into the Clipboard (usually stored in RAM) or the Scrapbook (stored on disk) with the intention of moving it, the information is saved in up to three forms:

- As a data file in the format specific to an application
- As an ASCII text file, without formatting
- As a picture file, in the QuickDraw format used by most Macintosh graphics programs

Not all programs or situations produce a Clipboard or Scrapbook in all three forms. In some cases involving large amounts of information, a program might give you a choice of how to store the Clipboard or Scrapbook.

The program from which you are planning to move information stores multiple formats because it cannot tell where the information is going; the safest strategy is to give the receiving program alternative formats to accept. When you choose to paste something from the Clipboard or Scrapbook, the receiving program checks the file formats and uses the most suitable format it can understand. Some programs may understand only their own Clipboard: You won't be able to paste information from another program into these programs. A few programs cannot read any Clipboards at all—which is sufficient reason to avoid them if you possibly can.

The Clipboard works like this: A text file could be converted into a picture of the text and merged into a graphics file, or a picture could be treated as a block and placed in a word-processing file to be printed along with the text. Because the original graphics and text files were created in such different ways, you won't usually be able to modify the picture file with the word processor or modify the text file with the graphics program. For example, MacWrite can use MacPaint information pasted in from the Clipboard or Scrapbook, and vice versa. MacWrite can also perform some simple changes to a MacPaint file included in text, such as moving or stretching it, but for any substantive changes, you have to go back to MacPaint.

How much memory is available at certain times imposes limits on the amount of information you can put in the Clipboard. You may have to move information in several blocks. Scrapbook size is limited only by available disk space.

The Clipboard remains intact if you switch disks, as long as you don't cut or copy something else into it. The Clipboard is also preserved between applications with MultiFinder.

The Clipboard and Scrapbook do have a major limitation: They can transfer static or "dead" information only. If you change the source information, you will have to transfer it again.

Live Links

With a live link, changes you make in one application are automatically made in all linked documents. Suppose you create a graph and incorporate it in a memo; with a live link, any change you make in the graph will appear immediately in the memo.

Live links are much easier to do within an integrated package than across separate programs, but no integrated package thus far permits live links in all directions. On the Macintosh, Modern Jazz from Lotus Development has the best live-linking features of any program, connecting spreadsheet and database to word processing. Microsoft Excel can set up live links among spreadsheets.

The Macintosh operating system contains a feature called Quick Switch Text (QST) that permits two independent programs to transfer live information, but because Apple has never actively supported this feature, QST has not been widely used. Microsoft Word, version 3, does use it, however, to transfer information from MacPaint, MacDraw, and Microsoft Excel.

Screen Dumps

If a program will not create a QuickDraw file for the Clipboard or Scrapbook, you can usually save a screen dump (the program's screen image) on disk by pressing Command-Shift-3. This operation creates a QuickDraw graphics file, called Screen0, on the disk. You can create up to nine screen files, named Screen1, Screen2, and so on through Screen9; at this point, you must rename earlier screens before you can continue. The screen dump is created by an FKEY (function key) program. (See Chapter 14.)

You can look at, modify, and print such screen dumps with MacPaint and other bit-mapped graphics programs. Within MacPaint, you can select the image and put part or all of it into the Clipboard. For the best printed results on a LaserWriter, however, avoid screen dumps if the program you're using can create a QuickDraw file; QuickDraw produces cleaner results.

If you don't need to modify the screen, you can print the active window by typing Command-Shift-4, or print the whole screen with Caps Lock-Command-Shift-4. This works with the ImageWriters, but Apple's software does not print screens on the LaserWriters. From user groups, you can get software (such as LaserKey) that adds LaserWriter screen dumps to Mac system software.

Reading Document Files

Nearly all programs let you move some information with the Clipboard and Scrapbook, but more complete information exchange requires that two programs be able to read the same document files. Document files come in many formats that have not been standardized, and, because programs are designed for widely divergent tasks, no single standard format is likely or possible. Some formats are shared by several different programs, however, and some programs can generate multiple document formats.

Most Macintosh programs store files in a document format unique to each program. Word processors store formatted text files, which include formatting information—typefaces and sizes, margins, spacing, and so on—along with text. But no two word processors create compatible document files, although several can read and write MacWrite files.

Other kinds of programs also create specialized document files. Spreadsheet programs use their own formats for the formulas and interrelations among spreadsheet cells. Object-oriented graphics programs, such as MacDraw and Microsoft Chart, store graphics in unique formats that are usually incompatible with other graphics programs, even though both incorporate elements of QuickDraw or PostScript. (See Chapter 10.)

You should not expect any application program to read another application's disk files, unless such an ability is specifically listed as a feature—particularly if the programs do completely different tasks. Although a collection of accounting programs might read a common file format, a word processor and an accounting program are unlikely to do so. A successful program—even one for a non-Mac computer—will attract other programs that can use its data. Many programs read MacPaint files, for example, and Microsoft Excel and other spreadsheet programs read Lotus 1-2-3 files.

When you have two programs that perform related tasks, however, file format compatibility is a critical requirement. Such compatibility can occur at several levels:

- Most desirable. Two programs can read each other's files directly. For example, if you use PageMaker, your word processor should be Microsoft Word or WriteNow because PageMaker can read and write files generated by both programs.
- Next most desirable. Two programs use a common file format that preserves all formulas and relationships. A Microsoft Excel spreadsheet, for example, can be saved in Lotus WKS format and read by Ragtime, a page-makeup program that includes spreadsheet features.
- Better than nothing. Two programs use a common file format that preserves current values or all text-formatting information, such as word-processing programs that exchange formatted files via DCA (document content architecture; see Chapter 21).

Special Links for Information Transfer

Where complete file compatibility is impractical, two programs that perform different tasks might have special provisions to pass along specific information. For example, when you construct a spreadsheet with Multiplan, you can copy values you want plotted to the Clipboard. Then, with Microsoft Chart, you can paste these values into the graph data. Chart lets you choose to paste data only once or to establish a link between files. With a link, you can change values in Multiplan, and Chart will automatically look at the same spreadsheet locations for new values to plot.

COPY PROTECTION

Software developers can take special steps to make their programs impossible to copy by ordinary means to protect themselves against people who copy and use the products without paying for them. Unfortunately, copy protection makes life more difficult for legitimate users, who clearly should enjoy both the security of backing up their programs and the freedom to put them on hard disks without going through cumbersome procedures each time they want to use the software.

The Macintoshes have no special built-in copy-protection hardware. Because effective copy protection often depends on hardware design, the introduction of the Mac II and changing disk drive hardware are making copy protection of Mac software more difficult. These problems, and strong objections from buyers, have led most major software publishers to drop copy protection in the past few years, except for games. Yet copy protection is unlikely to disappear completely.

Practical Choices

Whenever practical, you should avoid copy-protected programs because all methods of copy protection cause trouble sooner or later. In order of best to worst for the user, copy-protection schemes include the following:

- Unprotected. This is the least constraining choice because it allows you to make backups freely.
- "Validated" copies. You can make one or two working copies, including one on a hard disk; the working copy is completely self-sufficient, but you cannot make any more copies from it. The better programs let you "unvalidate" as well, so you can move a working copy if you change hard disks, for example. The validation process installs a hidden file or some other anomaly on the hard disk; the program looks for this quirk when starting up. Unfortunately, these anomalies pose a serious problem by interfering with routine backups of a hard disk. To avoid corrupting your hard disk, you might not want to install a validated copy of the program on it. You can instead simply copy the program in the usual way and then use a key disk (see below) to run it.

- Key information read once. This scheme requires you to insert a copy-protected "key" floppy disk when you start the program but not again as long as the computer stays on, even if you quit and restart the program. Aside from the initial brief check for the key information, the program can run from a working copy on another floppy or a hard disk. Some programs might require the key disk only once a week or every twentieth time you run the program.
- Key information read every time. This scheme requires you to insert the key disk every time you start the program.
- Key information read randomly. A few programs ask for the key disk at random times. This is disruptive because it's unpredictable; you may have to swap disks or interrupt your work to go hunting for the key disk.
- Key disk used to load and run the program. This is the worst-case situation, forcing you to run the program from the copy-protected floppy disk. A hard disk can serve only as a data disk.

The response to copy protection has been a proliferation of utility programs that can break the schemes. Advanced copy-protection-breaking schemes modify the program itself to bypass the copy-protection-checking routines, enabling the user to run programs from a hard disk without key disks or validated copies. As long as some software is copy protected, the battle between the designers of key disks and the authors of copy-protection-breaking programs will continue.

Ethical Choices

The debates that rage over software theft are part of the much larger question of copyright and the "ownership" of intellectual property. Software developers, like makers of sound recordings and book publishers, argue that both their creative and their manufacturing know-how have gone into making their product, and they should receive just compensation for it. Buyers argue that "just compensation" translates into exorbitant prices. At such prices, they reason, why not copy software put out by large, successful companies?

But price isn't the primary issue; low-cost software is probably stolen as often as the expensive variety. Small companies and garage shops are hurt as much as large companies and perhaps more. Some small companies and individuals distribute their products as "shareware": Their software can be legitimately copied, but a modest payment—a contribution, in essence—is requested from users. Even here, few users pay the requested fee.

With new technologies for copying audio or video recordings, printed matter, and software so widespread, what constitutes violation of copyright legally—or ethically—is becoming increasingly complex. But copying copyrighted material is hardly justified on the grounds that everybody does it. You always have one choice: If you don't like a software package because it has poor features, poor support, a high price, or copy protection, you can ignore it. The software industry is driven by the market; vendors will get the message.

MULTITASKING

Most of the time when you are working on your Mac, the CPU—which can process your commands much faster than you can issue them—is waiting for you to do something. One way to tap this unused power is via multitasking, whereby the CPU does more than one thing at a time and thus allows you to do the same.

In multitasking, one task—the active, or foreground, task—responds to the keyboard. Other tasks can be running in the background, but they do not respond to the keyboard until called into the foreground. On the Mac, both foreground and background tasks can have windows open on the screen, but only the foreground window is active.

Unlike their counterparts on mainframes and minicomputers, microcomputer operating systems were designed for single-user machines; consequently, true multitasking has been difficult to achieve on micros. On the Mac, multitasking could follow a number of schemes:

- Multiple loading (also called context switching). Two or more applications are loaded into memory at the same time, but only one—the foreground application—commands the CPU's full attention; background applications are frozen. You can instantly switch a background application into the foreground, however, because it is already in memory; you need not wait for a disk drive. Such a switch puts the foreground application into the background. Switcher (Finder version 4.0 and later), MultiFinder (shipped with Finder 6.0 and System 4.2), and Servant are utility programs that allow multiple loading. This is not true multitasking because background processes, except print spooling and communication, are not active.
- Round-robin processor control. After multiple tasks are loaded into memory, the foreground task commands the CPU's full attention; it then hands control to a background task, which has the CPU's attention until it hands control to another background task (or back to the foreground task), and so on. This round-robin procedure requires application programs to be specially modified.
- True multitasking. The CPU pays attention to all tasks, background as well as foreground, switching rapidly from one to another in a procedure known as time slicing. Background programs continue to operate in the background; their windows are not active during the time-slicing cycles. In simple systems, the CPU allocates identical periods of time to each task; in more sophisticated systems, time is allocated according to need or to priorities set by the user. Any single task runs more slowly with time slicing than without, although the degradation may not be perceptible in ordinary applications. Real-time

operation is a kind of true multitasking in which one process that urgently needs computing power can grab all the resources it needs without waiting.

True multitasking is hard to do; the more limited multiple loading is much easier to implement on Macintoshes. In practical terms, the advantages of true multitasking might be obscure on single-user microcomputers. Although it may be handy to have your database program sort files in the background, a spreadsheet waiting in the background usually has nothing to do.

NON-MACINTOSH OPERATING SYSTEMS

As a Mac user, you might never need to use another operating system, but some systems that were originally designed for other computers can run on a Macintosh, particularly on the Mac II.

Because they were designed for other CPU chips, most non-Mac operating systems do not run directly on a Mac because each chip type follows a different set of instructions. For example, the operating systems for the Apple II run on MOS Technology's 6502 chip; TRS/DOS and CP/M-80 run on Zilog's Z80 chip; and MS-DOS runs on Intel's 8088, 8086, 80286, and 80386 chips. Many microcomputers besides the Macintosh family use Motorola's 68000 and 68020 chips, but these computers have operating systems that generally require a different arrangement of hardware from the Mac's.

If you run another operating system on a Mac, the interface will most likely be nonvisual. For this reason alone, you will probably not want to use other operating systems unless:

- You have a large existing inventory of useful programs that you cannot afford to replace.
- You have written many programs in a form not transferable to the Macintosh.
- A program you use includes accessories that Macintosh software does not support.
- You need to work directly with files from another computer system.

Of the various non-Macintosh operating systems that can run on a Mac, two stand out: MS-DOS and UNIX.

MS-DOS on the Mac

MS-DOS (Microsoft Disk Operating System) is the standard operating system for IBM microcomputers and the many clone computers that mimic IBM's models. It is relatively simple, with a cryptic command-line interface. PC-DOS is what IBM calls the version of MS-DOS that it sells; among IBM PC users, the name is often shortened to DOS.

Because the IBM PC is so popular, a great variety of software runs under MS-DOS. Most of these programs are not as good as their Macintosh equivalents, although there are more highly specialized programs for the PC than for the Mac. If you are a Mac user who needs to deal with the MS-DOS world, you have several choices:

- Working with MS-DOS data. If you need to work with data generated by an MS-DOS program, you may or may not need the program itself. Microsoft Excel on the Mac, for example, can read and write files generated by the MS-DOS spreadsheet program Lotus 1-2-3. In such cases, you will need only a way to move data stored on an MS-DOS disk to a Mac disk—something that can be done by a variety of means. (See Chapter 21.) If, however, the MS-DOS program stores data in a format that no Macintosh program can understand, you will need to run the MS-DOS program to get to the data.
- MS-DOS software emulation. MS-DOS can be run on a Macintosh II with software from Insignia Solutions (San Francisco, CA) that makes the Motorola 68020 CPU behave like an Intel 8086 CPU chip; no additional hardware is needed. Because emulation requires that many processor instructions be translated from one chip to the other, it almost always means sluggish performance, but at least an MS-DOS program will run.
- MS-DOS boards. An MS-DOS processor board that contains an IBM PC processor (either Intel's 8088, 8086, 80286, or 80386 chip) can be added to the Mac Plus, Mac SE, or Mac II so that MS-DOS software can run; such boards cannot accept IBM hardware accessories, however. Video performance is usually slow because the IBM-style display must be converted into a Macintosh window.

  PerfecTek (Milpitas, CA) has an 8086 board that clips onto the 68000 processor in a Mac Plus.

  AST Research (Irvine, CA) makes two MS-DOS boards, the Mac 86 for the Mac SE and the Mac 286 for the Mac II. The Mac 86 board plugs into the SE's internal expansion slot and contains an Intel 8086 processor that runs at an effective 4.77 MHz, the same speed as the original IBM PC. The 8086 processor shares memory with the Mac SE's 68000 CPU. The more advanced Mac 286 product mimics an IBM PC/AT. It consists of two boards that plug into a Mac II: One has an Intel 80286 processor running at 8 MHz, the second contains memory chips that are separate from the Mac's own memory. The Mac II's 68020 CPU handles input/output and the video screen; the Mac 286 system includes disk-controller hardware for connecting an external 5¼-inch floppy disk drive. The system software supplied with the Mac 286 allocates space on the Mac's hard disk drive for storing MS-DOS files; these files can be read by Macintosh applications.

  An accessory board for the Mac II with an 80386 chip may also appear by the end of 1988.
- External MS-DOS hardware. MacCharlie, made by Dayna Communications (Salt Lake City, UT), is an MS-DOS computer built into a case that fits around the small Macs. MacCharlie contains its own 5¼-inch disk drives but uses the Mac keyboard and screen.

- Buying a standard IBM PC or clone. A complete IBM PC clone, including disk drives and video display, costs about as much as the equivalent MS-DOS hardware accessories for the Macintosh, which are simply not made in large enough quantities to justify low prices. Whether it is IBM-made or not, a separate MS-DOS computer will give you both software and hardware compatibility, and you will be able to use the many accessory boards made to fit inside an IBM PC. And you can still move information back and forth between the PC and the Mac. (See Chapter 21.) The main disadvantage is size: Two desktop computers take up much more space than a self-contained Mac equipped with an MS-DOS board.

  If you buy an IBM-compatible computer to run MS-DOS programs, you should get an 80286-based model (IBM PC/AT, PS/2 models 50, 60, and compatibles) or an 80386-based model (COMPAQ 386, IBM PS/2 model 80, and compatibles), not the obsolete 8088/8086 designs. Both 80286 and 80386 computers can run versions of OS/2, a new operating system with more sophisticated features than MS-DOS.

UNIX

UNIX, a multiuser, multitasking operating system, was originally developed for minicomputers by Bell Laboratories in the early 1970s. Widely adopted by academic institutions, research laboratories, engineering enterprises, and government agencies, UNIX has evolved over the years into a family of operating systems with many variants. The two major branches of UNIX are AT&T System V and Berkeley 4.2, which has features developed at the University of California. Each of these branches has versions for different computer systems, most often minicomputers and expensive engineering workstations.

UNIX is a powerful operating system whose advantages and disadvantages, particularly for microcomputer users, have been hotly debated. The UNIX system boasts many utilities, which programmers often swear by, but its command-line interface is so cryptic that several replacement interfaces, or shells, have been developed to make it easier to use. UNIX is enormous: Unlike the Mac operating system, which takes up only a fraction of a single floppy disk, UNIX requires dozens of floppy disks—so many that it is often shipped on a hard disk or on a 20 MB or larger tape cartridge.

When computer processors were expensive, and local area networks did not exist, UNIX was considered ideal for places like universities or large research facilities because with enough video terminals and keyboards, many users could share the same CPU. Now that processors are cheap and networks are increasingly common, multiuser operation is not so clear an advantage.

Because UNIX runs on a variety of CPUs, it has been considered highly portable, permitting the same program to run on many different computers, even those with significantly different hardware. In reality, however, this has proven only partly true; virtually all software requires modifying for each UNIX computer and each UNIX variant. For this reason, and because the operating system is so clumsy to use, most institutions with a significant investment in UNIX products have a resident UNIX guru whose primary mission is solving the many arcane problems that arise.

So what if you need or want to run UNIX on a Macintosh? UNIX can run on the Motorola 68000 chip, but the hardware configuration of the small Macs cannot cope effectively with it. A powerful workstation in its own right, the Mac II is more suitable; indeed, Apple is distributing A/UX (Apple/UniX), which was developed in conjunction with UNIX developer UniSoft.

A/UX is based on AT&T's System V, version 2, release 2, with many Berkeley 4.2 utilities included; also supplied are the Bourne, Korn, and C command shells, which install menus that hide the standard interface. A/UX can work with AppleTalk, Sun's Network File System, and B-Net, which is an EtherNet-based protocol following the Department of Defense's TCP/IP specifications. To run A/UX on the Mac II, Apple recommends that you have at least 2 MB of RAM and an 80 MB hard disk.

For its multitasking functions, A/UX requires hardware memory management, a circuit that gives each task its own memory area to work with; thus, even if one task crashes, the others won't be affected. For A/UX, memory management is done by a Motorola 68851 chip installed in the Mac II. (This chip should be installed only if you are using UNIX because it slows normal Mac operation.) A single Mac II can be used as a UNIX host for multiuser operation as well. Each individual terminal should be connected to a Mac II through special serial ports such as those on AST's ICP (intelligent communications processor) input/output board.

Unlike the Macintosh operating system, UNIX is text based; its core software does not process bit maps. Because many UNIX computers have no graphics capabilities, most UNIX application software also shuns graphics. One class of modern graphics software does work under UNIX, however: computer-aided graphics and design, page-makeup, and other graphics programs that run on engineering workstations made by companies such as Sun, Apollo, and DEC. Because UNIX has no graphics standard, software developers must not only adapt their programs to the UNIX variant specific to a particular workstation model, but they must also tailor the graphics portions of their programs to each workstation.

Many of these programs are highly sophisticated, and will undoubtedly be brought over to the Mac II. They could appear as:

- UNIX programs with their own graphic interfaces specifically adapted to Macintosh hardware.
- UNIX programs that rely on the Mac ROM to create windows, menus, and graphics.
- True Macintosh programs running under the Mac operating system and using the Mac ROM and interface.

The X Consortium, a group based at the Massachusetts Institute of Technology, is now trying to establish a graphics interface for UNIX (and some other operating systems) called X Windows. Corporate members of the consortium include Apple, AT&T, IBM, Sun, DEC, Sony, Hewlett-Packard, and others. In its initial release of A/UX, Apple does not support X Windows or any graphics format used by earlier UNIX computers; to run on the Mac II, UNIX software that requires graphics on screen must therefore be modified.

A/UX does, however, support a variety of text-based video displays. If the Mac II is used as a multiuser UNIX computer, the terminals connected to it will probably handle text only; sending text to the terminals is much easier than sending bit-mapped graphics.

Other Operating Systems

Without additional hardware, a Macintosh can run several operating systems designed for earlier microcomputers, although the application software available for these operating systems is badly outdated.

- Apple II programs can be run with software, such as II in a Mac from Computer Applications (Raleigh, NC), that emulates the Apple II's 6502 CPU chip. The II in a Mac program includes communication software to move programs and data from an Apple II's 5¼-inch disk drive (copy-protected Apple II programs cannot be converted unless they are modified); it can read 3½-inch Apple II Unidisks directly.
- The UCSD (University of California, San Diego) p-System is available for the Mac from Pecan Software (Brooklyn, NY). The p-System was an early effort to create a single operating system standard so that the same program could run on a variety of computers. Unfortunately, this universal approach led to sluggish performance and so many compromises that it has been mostly a failure.
- Mac versions of several other operating systems have come and gone. At least three companies marketed CP/M-80 products for the Mac, and one company did a version of CP/M-68K; all seem to have disappeared.

17: Problem Solving and Getting Help

Your Macintosh will need only minimal maintenance and probably few, if any, repairs. If you run into a problem you suspect is due to a hardware or software failure, try reproducing it on another Macintosh (the same model, with the same version of the System and Finder as yours). If the problem is reproducible, the failure lies either in the software or in the disk itself; try a different disk (but not a copy of the suspect disk). If the problem persists, you probably have a software bug; if not, the first disk was most likely faulty.

If the problem cannot be reproduced on another Macintosh, your hardware may be responsible. For difficulties not discussed below, and for more information, check your manuals, consult your dealer, or contact a user group.

FAULTY CABLES

On any computer, the most failure-prone parts are the plugs and connectors. If you are having difficulties with an external accessory, such as a printer, look at the plugs carefully. Are any pins bent? Some pins may be missing—which is normal—but if you're not sure if the right ones are missing, turn off the power to both the Mac and the peripheral, and try swapping cables.

If you don't want to turn off the power, thereby erasing everything in RAM, you can usually safely unplug and check the keyboard, mouse, and serial port cables with the power on; follow the manufacturer's recommendations. Always turn off the power before unplugging anything connected to a SCSI port or a floppy disk port and before disconnecting accessory boards plugged into Mac SE or Mac II expansion slots.

GLITCHES

All computers are subject to glitches that cause temporary setbacks but might not be symptoms of anything serious. Static electricity can sometimes cause difficulties. Many companies sell antistatic pads to put under your chair or to stand on, or pieces of metal you're supposed to touch before operating the computer. These products are effective only if grounded properly; they are usually unnecessary unless the air is so dry that you are uncomfortable and you set off sparks every time you touch a doorknob.

Any computer will occasionally seize up so that you can't get it to respond. With a Mac, this kind of glitch can be disconcerting because you may still be able to move the mouse pointer, which uses a portion of memory separate from application programs. Sometimes you won't be able to pin down the exact cause. If a software bug is the culprit, the computer should seize up again if you repeat exactly each step that preceded the lockup. But you won't always remember the exact steps; besides, something you did hours earlier could be the cause. There could be incompatibilities among your application program, desk accessories, and FKEYs; the problem might be the power line. A bit of dust can make a disk temporarily unreadable: After the disk drive head has gone on to read another part of a floppy disk, it may be able to come back and read a track that was faulty minutes earlier.

Some problems simply have no explanation. You may not be able to repeat some glitches, and you will never know why they occurred. Memory chips are susceptible to errors resulting from bit changes induced by cosmic rays, although such errors are rare—far fewer than one a year. Rather than spend hours trying to figure out a problem that you cannot replicate, accept that these things happen from time to time, even with the largest, most expensive computers.

If you work with programs that often lock up your Mac, install the plastic piece labeled INTERRUPT RESET on the vent grille. (See your Macintosh manual.) Press RESET; the Mac will restart as if you had just turned it on. Once this piece is installed, take care not to have anything push against it, or you will unexpectedly reset the computer. (The INTERRUPT button is intended for use with a debugging program in software development.)

When all else fails, you may have to turn off the power to the computer and peripherals and go away for awhile; the problem may have disappeared by the time you return. Do this only as a last resort, however. If you have a hard disk drive, be sure to leave the power off for several seconds—long enough for the hard disk drive to stop spinning. If the power is not off long enough, some equipment will not go through a power-on procedure (the capacitors in the power supply continue momentarily to furnish power even after the main switch is off), and your problem may persist.

To turn off a locked-up Mac II, press the power switch on the main case once (not the Power-on key on the keyboard).

IF A FLOPPY DISK GETS STUCK

If your computer locks up, you can eject a floppy disk in one of three ways (two of which will erase RAM, and all of which could make the disk, or parts of it, unreadable):

- Press Command-Shift-1 or Command-Shift-2 to force the external or internal disk to eject.
- Turn off the power, press and hold down the mouse button, and turn the power back on. The disk drive will eject the disk in a moment.
- If the power has failed, remove the disk mechanically by inserting a heavy pin (a heavy paper clip bent open) into the small hole just to the right of the disk drive slot.

If the power has failed, you can remove a floppy disk with a heavy paper clip.

POWER-LINE PROBLEMS

Many companies sell power-line conditioners, costing $8 to $100, that eliminate short, sharp voltage pulses or the high-frequency interference that can induce errors in a computer (they rarely cause lasting physical damage). Whether you need a power-line conditioner depends on the quality of your power source and whatever else is connected to it. If your Mac locks up or produces errors whenever an appliance is turned on in the building, try a power-line conditioner.

If you have a choice, do not plug your Mac into an electrical circuit that is used for any heavy equipment, such as air conditioners, refrigerators, power tools, photocopying machines, or vacuum cleaners. The motors in these devices can generate a transient power-line spike that occasionally spells trouble. Vacuum cleaners are especially insidious because they move around, unlike refrigerators. If you see a vacuum cleaner coming your way, save your file on disk and then continue working; you probably won't have a problem, but it's best not to take chances. Once you're sure a potential hazard causes no trouble, you can ignore it.

Brownouts

A Macintosh is fairly resistant to brownouts, or low-voltage periods. The Mac Plus, for example, is specified to operate on 105 to 130 volts, but it actually runs satisfactorily on as little as 95 volts. (Some units even run on 75 volts.) The international version has an even more tolerant power supply, specified for 90 to 130 volts and switchable to 180 to 260 volts. This version will actually run on anywhere from 85 to 135 volts and 170 to 270 volts. The Mac SE and the Mac II have similar voltage tolerances, although any hard disk drive installed in the computers may be more sensitive. The Mac II has a power supply that can switch automatically between 120 and 220 volts.

The power-line frequency doesn't matter; Macintosh specifications call for 47 to 64 Hz, but the power supply will actually run on a much wider range. Leakage current increases with higher power-line frequencies, so be sure your computer is properly grounded. The built-in clock/calendar uses its own quartz oscillator and battery; it does not depend on the power-line frequency.

Severe Electrical Storms

During severe thunderstorms, unplug the computer and disconnect the modem from the telephone line.

Power Failures

Power failures are another threat. Whenever the power fails, everything in RAM disappears, and any open file on disk can be lost as well. There are several steps you can take to combat this threat; three require purchasing more accessories, and the fourth is simply a question of work habits.

Uninterruptible and standby power supplies

You can buy a standby or an uninterruptible power supply. Both are built around a battery, a battery charger, and an inverter (which converts the battery's DC output to AC). If your computer is connected to a truly uninterruptible power supply (UPS), it is always running off a battery that is continuously recharged, so power to the computer stays constant regardless of what is happening to the AC power. Such a system is fairly expensive, for its circuitry must be well enough designed to regulate charging and discharging without damaging the battery.

With a cheaper standby power supply (SPS), the computer normally operates directly off AC power; if the AC power fails, however, the unit switches to battery operation. The switch happens so fast—generally in less than 15 milliseconds—that the computer does not notice the changeover. (The capacitors in the Mac's own power supply can keep its logic circuitry alive for more than 50 milliseconds.)

The battery in a UPS or SPS system usually stores enough power to last 15 to 20 minutes—plenty of time for you to close all your files and turn off the machine gracefully, *if* you are present when the power fails; some UPS and SPS units have much larger batteries. Magic Software (Bellevue, NE) offers a desk accessory called AutoSave DA, which periodically saves any file you are working on; with this automatic feature, you lose only a few minutes' work even if the power fails while you are out to lunch.

For a small Mac, a UPS or SPS unit needs to supply at least 100 watts of power (150 watts if you have a hard disk); for a Mac II, 350 watts will cover even a color monitor. UPS and SPS systems normally include power-line conditioning. If you install a UPS or SPS, peripherals such as printers need not be plugged into it; protect only critical components—the Mac itself, hard disk drives, and the video monitor.

Permanent battery operation

Continuous battery operation may be advisable if you are working in an area with exceptionally unreliable power, or if you have no conventional power source. Under such conditions, the small Macs are more practical than the Mac II because they require less power. When you aren't running your Macintosh, the battery can be charged from an AC power supply or generator. An inverter with a square-wave power output works satisfactorily (although in rare cases an isolation transformer may be necessary to smooth the square wave slightly). But because batteries require considerable maintenance, avoid using them unless you have to.

Recovering disk files

A few programs automatically save information in RAM on disk whenever you haven't entered anything from the keyboard for some time. Such automatic disk writing is a valuable feature if it means you can recover the saved information after a power failure. Many database programs include automatic disk writing, and AutoSave DA can install the feature in a variety of programs.

Disk-recovery programs such as 1st Aid Kit (1st Aid Software, Boston, MA) examine a disk and recover any files or parts thereof that can still be read, even if the disk is unreadable by normal means. Sector readers are programs that read raw data from a disk sector by sector and display them in ASCII, binary, or hexadecimal (base 16) form. A file is not necessarily written on contiguous sectors of a disk; to recover an entire file with a sector reader, you must determine which sectors on the disk contain that file. Sector readers allow you to change any byte on a disk, which makes them a handy tool for experienced users; however, if you change a byte in the wrong place (in the part of the disk that lists the locations of files, for example), you can make your disk totally unreadable. The programs Fedit (MacMaster Systems, Sunnyvale, CA), MacZap (Micro Analyst, Austin, TX), and MacTools (Central Point Software, Portland, OR) all include sector readers.

For a thorough discussion of hard disk problems, see the article "Hard Disk Medic" by Bob Johnson in *MacUser*, December 1987 (adapted from Macintosh Technical note 134, published by Apple Computer).

Cautious work habits

One of the most undramatic and yet surefire ways to ward off the threat of power failures is to habitually safeguard your work.

Saving frequently on disk is one excellent habit to get into. If there is a power failure, you will generally lose only whatever work you've done since you last saved. So be prudent and save after every complex operation, especially when moving information between programs.

Printing backups is a simple mechanical way of preserving some types of work. Although this is not ideal because everything needs to be retyped, having a printed backup is easier than having to rethink it all.

If the program itself is faulty and won't print your document, you can sometimes still use the program to read the document and copy everything to the Scrapbook; then at least you have a record that you can print from another program.

Backing up your disks is another good habit to get into. Copy important data regularly, and make at least one backup of the programs you use. (See Chapter 6.)

MAINTENANCE

Although the Macintosh is not like an automobile that requires periodic tune-ups to stay in shape, here are some maintenance tips.

Routine Cleaning

If you don't use a dustcover, dust will eventually build up inside the Mac. Modest amounts won't affect operation. With or without a dustcover, the exterior will probably need occasional cleaning. Use a damp cloth and a small amount of a mild, nonabrasive household cleaner if necessary; take care not to get anything dripping wet. Clean the screen with glass cleaner; for optical glass antiglare screens, use lens cleaners only. (See Chapter 4; Chapter 5 tells you how to clean a mouse.)

There is little agreement about cleaning floppy disk drive heads. Many microcomputer users have never cleaned their disk drive heads despite years of heavy use and have had no problems; others claim that regular cleaning is essential. Certainly you should not have to clean the heads more than once or twice a year under normal conditions. (If you have a hard disk drive, your floppy disk drive won't get dirty quickly because it won't be used much.)

Long-Term Maintenance

With long-term use, the Macintosh's mechanical parts may require service and alignment. The floppy disk drives and the printer may need service after several years. A hard disk drive may eventually require new bearings and other adjustments, but when that happens, you may find that getting a new, higher-performance hard disk will be a better investment than repairing an old one. An aging CRT will begin taking longer to start up, and the screen will get a little dimmer. The rest of the computer should never require service unless a part fails.

The internal designs of the Macintoshes make repairing them very difficult for casual tinkerers. Most of the chips are soldered in, and the boards have four-layer circuits; the chance of damage is very high if you attempt to unsolder components with an ordinary soldering iron.

Service Contracts

Vendors like service contracts because they are a high-profit item. But for users, service contracts rarely make sense; microcomputers are relatively cheap and do not often break down. If your company has multiple microcomputers, you should buy a spare unit instead of multiple service contracts. If something breaks, simply swap computers. Realistically, though, the spare computer almost always winds up in constant use, so you have to get another spare....

PHYSICAL PROTECTION

Aside from a dustcover, a Macintosh in normal usage needs little, if any, protection from the elements. Water from a sprinkler system or foams from fire extinguishers will damage any electronic equipment. The fire-fighting gases Halon 1301 or Halon 1211 can put out fires without such damage.

If theft is a problem in your area, several companies make locking devices for the major components.

GETTING HELP

Almost every computer user needs help at one time or another. Whether you're trying to find out what you did to make that crucial document disappear or how to connect some esoteric accessory, there are resources that can help you.

Manuals

Your first source of information should be the manuals that come with the Macintosh or with the hardware or software that's giving you a problem. Apple's manuals are among the best in the microcomputer business and bear careful reading.

Unfortunately, however, manuals have never been a strong point of the microcomputer industry; good ones are so rare that a case can be made for selecting software on the basis of the manual's quality alone. Generally, a company that puts the effort into a good manual is also a company that cares about its customers.

Hotlines

Many microcomputer companies maintain hotlines staffed by people who are trained to answer technical questions directly. Others do not have hotlines, insisting instead that you direct queries to your dealer, and accepting technical questions only from dealers. Through such a procedure, your questions will often be garbled or misunderstood. When possible, buy only products supported by a hotline.

Remember, though, that hotlines are expensive to staff and maintain, so before calling a company's hotline, make a reasonable effort to find the answer to your question in your manuals. Don't expect answers to questions about products other than those produced by the company you are calling. If you are calling because a program failed ("crashed"), be specific about exactly what happened, any error messages you might have seen, and any action you might have taken that could have caused the crash. If you are using any nonstandard component—hardware or software—note that as well.

For some products, the company may charge for hotline access. Find out whether competing products have such charges and compare carefully.

Computer Stores

You can, of course, always go back to your dealer, but dealer competence varies widely. Some dealers are helpful far beyond the call of duty; others are unable to answer even the most straightforward questions. Generally, salespeople should be able to answer simple queries and check the operation of your unit against others. Beyond that, it's hard to tell you what to expect.

If possible, ask your friends where they bought their Macs, accessories, or software and whether they are satisfied with dealer support. If your friends are pleased, then you should buy from that dealer also. Avoid dealers with a poor reputation for service, even if their prices are a little lower. Computers are complex enough that most people will need help from their dealer at some time.

As with any other fast-growing industry, the microcomputer business has attracted fly-by-night operators. Check with user groups if you have doubts about a dealer; occasionally a Better Business Bureau might have useful information. Always be cautious if you are asked to pay in advance for any item not immediately available. For an unusual item, a small cash deposit may be reasonable, but you should never have to put a deposit on anything a store regularly carries.

A salesperson who seems unusually determined to sell you a particular product may be getting a "spiff," a direct sales commission from the manufacturer, above and beyond the normal dealer discount. Spiffs distort the sales process, but, unfortunately, they are common.

Mail-Order Companies

Computer magazines are full of ads from mail-order houses, generally promising quick service at low prices. Doing business with these outlets can be a good deal or a disaster. Many mail-order companies accept your order, charge your credit card, and do nothing for as long as possible. The honorable ones ship your order when promised and submit the charge only after shipping the product. It's hard to tell the two apart; word of mouth helps, but a company's quality can change over time.

Most mail-order houses offer little or no support for the products they sell. Once you've bought something, they don't want to hear from you until you place another order. By eliminating support, they can sell for less than normal dealers.

There are a few good mail-order houses that offer excellent support by telephone—much better, in fact, than most dealers. If you find one, you're lucky. Tell your friends (but don't be surprised if the quality of support withers as the company grows).

Whatever you do, don't take advantage of a local dealer to get information and then place your order with a mail-order house. Besides being unfair, doing so discourages the competent stores that deserve a return on their efforts. A dealer offering useful presale information will also provide useful service after the sale.

Publications

Magazines and newsletters about the Macintosh started quickly. Like everything else, the quality of information in publications varies. The computer press grew so rapidly that desperate publishers hired many naive writers and editors. The best magazines are staffed by technically sophisticated journalists; they bring you the latest information along with a little hype.

Weaker magazines can be identified by their style: articles that consist mostly of quotes from computer-store sales personnel—genial, uncritical, and uninformative. A few magazines have gone so far as to ban all criticism of products. Paradoxically, some of the worst magazines make the fewest errors because they offer so little information in the first place.

The quality of computer books varies too. Most Mac books now concentrate on a single software product. The value of these books depends in part on the quality of the manual that comes with the software. Books inevitably take longer to produce than magazines, so you shouldn't look to books—including this one—for the latest information. Look to them instead for insight or as a reference. Good books anticipate questions and problems, and if they can't always solve a problem, they can give you tools for finding your own solution.

Introductory books about the Mac are no longer common, partly because Apple's manuals are so informative. The shelf life of books describing shortcuts and tricks that speed up work for experienced users is also short because new versions of software and hardware appear so frequently. *The Macintosh Bible* (Goldstein & Blair, Berkeley, CA) addresses this problem by offering updates with the purchase of the book.

User Groups

You can greatly expand your circle of Macintosh-user acquaintances and increase your source of word-of-mouth recommendations by joining a user group. To locate a group in your area, contact a local Apple dealer, or check the notices in a local computer newspaper—or you and your friends can start your own. Two large user groups have a national following: The Boston Computer Society (1 Center Plaza, Boston, MA 02108) and the Berkeley Macintosh User Group (1442A Walnut St. #62, Berkeley, CA 94709).

A good user group is probably the best single source of information about computers. Club members run the gamut from rank beginners to computer engineers, and everyone shares information. If you can't get a direct answer to a question, you will probably find someone with a similar problem, and together you can find the answer.

Often the most sophisticated users are "hackers," people who spend every spare moment bent over their computers, probing some obscure hardware or software feature or bug. Such enthusiasts are like hot-rodders: They may be a good source of information about repairing your carburetor, but their ideas about how to choose and use cars may have little in common with your needs. Yet the best hackers will understand the computer's innards far more thoroughly than any dealer, and they can be an invaluable information source.

If you are thinking about buying a Mac, you can get advice from owners of competing brands. Many people who spend several thousand dollars on a computer system develop an emotional attachment to their investment and lose their objectivity. If you buy a Mac, you may fall prey to the same syndrome. Keep in mind that no microcomputer is perfect; all designs are compromises. Most major current models do at least a passable job for average tasks. Every computer, including yours (whatever model), will be obsolete someday.

Electronic Bulletin Boards

Electronic bulletin boards can be another source of information. These bulletin boards are stored inside a central computer, and you gain access to them by modem and telephone lines. The simpler bulletin boards run by user groups are usually free. Macintosh bulletin boards on the national electronic database services run by CompuServe and Delphi cost $5 to $20 per hour and don't necessarily offer any more information than a good bulletin board run by a user group. Bulletin boards can be effective for short, clearly defined questions, tips on many subjects, and getting public-domain and shareware programs.

Electronic bulletin boards have generally been designed for hobbyists, and none of them, including commercial services, work well. The Macintosh interface is improving matters, however, and products such as CompuServe Navigator make access to the information in CompuServe much easier by replacing a slow menu-driven interface with a graphical one.

Sectio

For More Information

Four

For the dedicated and the curious, this final section offers a potpourri of Macintosh facts and ideas—more detail on topics brought up in previous chapters, plus some new material. Some of the presentations are highly technical, written for specialists and intended primarily for reference.

18: A Macintosh Medley

Since personal computers first appeared on the market more than a decade ago, they have become indispensable in nearly every business and profession. Although the early models were heavy and unwieldy and were limited to games and a few "serious" applications, microcomputers now can perform hundreds of unforeseen tasks—including many that extend the reach of the disabled and the traveling range of the once officebound. This chapter surveys some of these uses, closing with a comparison of the Macintosh and IBM's personal computers.

WITH SOME EXTRA HARDWARE...

With the appropriate hardware accessories, a Mac can adopt many roles; a sampler of accessories follows.

Telephone Management

Several companies are developing telephone products for the Mac to assist voice calls rather than transfer data. One system will be a smart telephone. The Macintosh will plug in between a telephone and telephone jack, and software on the Mac will dial, time, and log phone calls. Disks—from floppy to optical—will store telephone directories of all sizes, putting a host of phone numbers at your fingertips.

If, in such a system, the Mac and telephone were connected to a PBX (private branch exchange), the software would display PBX commands on screen—whom to call, where to transfer a call, and so on; a hardware interface would send appropriate instructions to the PBX.

A Macintosh could also replace the operator's console in a PBX system, where its visual interface would make operating the PBX far easier than with traditional consoles. Similarly, a Mac could be used to program a PBX for where a line will ring, which lines can call out, and so on.

Video Control

Videodisc- and videotape-controller hardware connected to a Mac will provide all the functions of random-access tape and disc players. The video equipment will generate a full-motion image that can be displayed on a standard television or monitor while the Mac screen continues to work as it normally does. On a Mac II with suitable hardware accessories, the computer and video images could even share the same screen. (See Chapter 21.)

A Macintosh will also control videotape recorders and film-editing equipment, with a hardware interface that reads and generates standard SMPTE (Society of Motion Picture and Television Engineers) time codes. A video digitizer will grab the frames you designate, creating miniature images on the screen. You will then be able to edit by manipulating the images. This is a great improvement over traditional systems, in which you must memorize the images. The software will automatically assemble the finished tape by issuing commands to the recorders and video-effects equipment.

Sound and Music

By itself, the Macintosh can synthesize musical sounds adequately for teaching and demonstration, but serious musicians will want to use it to control professional musical instruments, keyboards, and synthesizers. Any Mac can be connected to the many electronic instruments that use the musical instrument digital interface (MIDI) for electronic controls and interconnections. MIDI is a 31.25-kilobits-per-second serial protocol, supporting 16 polyphonic channels on each circuit. A dozen software packages that work with MIDI interfaces are already available for Macs. MIDI interface hardware is available from Apple and independent vendors; it plugs into the printer or the modem port.

Accessory boards for the Mac II will take on much more ambitious sound processing, including sound synthesis with full 16-bit quality—the same as compact discs. With a high-speed hard disk controller, such a system could record, edit, and play back with the same technical capabilities as a small recording studio. Since sounds can be manipulated with software, many effects could be created without additional hardware.

A Macintosh will also control multitrack tape recorders and digital mixing boards in recording studios, marking and displaying sound fragments as a compressed spectrogram; as with video accessories, graphical editing will be possible.

Instrument Control

A wide variety of devices (thermometers, machinery, fuel-tank gauges, and so on) can be coupled to a computer for automated data collection; an instrument controller converts the signals from the data-collecting device into a form understandable by the computer. Analog-to-digital (A-to-D) converters change analog voltages into digital signals, and digital-to-analog converters turn digital signals into analog voltages. These devices are designed mainly for laboratory or manufacturing applications.

Instruments can be connected to the computer via several routes, depending on how often you need to monitor (or sample) the instrument and how much information is in each sample. The slowest connection is through a serial port; a SCSI port is much faster, and an accessory card that plugs directly into the bus in a Mac II is even faster.

Many companies make controllers that connect through RS-232C serial ports. A serial port controller can also manage instruments that operate on the IEEE 488 bus, which can handle up to 15 devices at a time. Metaresearch (Portland, OR) sells software and hardware developed by Reed College, including an icon-driven interface for controlling an instrument network where each device has an Intel 8031 processor.

Once acquired, the data from instruments can be manipulated and analyzed on the Macintosh with software written in virtually any programming language. To save you the trouble of writing your own programs, most instrument manufacturers either sell or can recommend software that extracts data and puts them into a spreadsheet or other application program for analysis.

A digital oscilloscope attachment can turn the Mac into an electronic tool for analyzing waveforms. The computer can be used to store and compare electronic signals. A logic analyzer lets you trace digital circuits. Signal generators use the screen to display waveforms; analyzers display the distortion products. An entire electronic testing set can be built around a Macintosh.

Power Controllers

Power controllers working with X-10 switches let you turn appliances on and off with your computer. Computer-operated controllers make sense for industrial applications and for the disabled; they are less useful for ordinary offices or households.

TRAVELING MACINTOSHES

Although you can carry the small Macintoshes more easily than conventional personal computers, they are not truly portable. They consume too much power to run off lightweight batteries, and you can't use one on your lap; a laptop Mac requires completely redesigning virtually every component. Independent companies cannot build a truly portable Macintosh because Apple will not license the right to replicate critical Mac circuits, including the Mac ROMs. Yet despite the refusal to grant licenses, Apple's own development of a portable Mac is proceeding slowly.

If you don't want to wait for Apple, several companies have adapted Macs for travel by buying computers from Apple and repackaging the circuits. Unfortunately, none of the early models is very satisfactory. Dynamac Computer Products (Golden, CO) has made the sleekest-looking model, although some of that sleekness results from omitting a place to store the mouse. Intelitec (Fairfield, IA) and Colby Systems (Fresno, CA) have also made laptop Macs, although their models are nearly as heavy as a Mac Plus.

In building their laptop Macs, all these companies face the same problems: Mac circuits are not designed for low power consumption or small size. Mice are a nuisance on any computer that is intended to be run on your lap. What's more, no compact flat-panel screen technology delivers both high legibility and low power consumption. (See Chapter 4.)

If you use your Mac in an airplane, you need to consider some potential problems. The Macintosh is certified to be sufficiently free of electromagnetic radiation for home use (FCC type B), but there are no certification procedures to qualify a device as safe to use on aircraft. Modified Macs with electroluminescent or plasma screens can interfere with the airplane's radio-based navigation system because these screens rapidly switch high voltages (100–200 volts) to turn pixels on and off; the higher the voltages switched, the more radiation emitted. Computers with LCD screens generally interfere much less because only a few volts are switched to drive those screens. Before you use your Macintosh on a private aircraft, consult the companies that supply the avionics to see if they have advice or tests concerning possible interference from the Mac. In any event, prudence suggests that you not use any device that might radiate interference during takeoff or landing, or while flying near potentially hostile airspace.

An unpressurized aircraft is one of the few places where a Macintosh might be used beyond its altitude limit of 10,000 feet, or 3000 meters. (See the Appendix for an explanation.)

THE INTERNATIONAL MACINTOSH

The Macintosh is available with system software in many languages: North American English, United Kingdom English, French, Canadian French, German, Spanish, Dutch, Italian, Flemish, Swedish, Norwegian, Portuguese, Greek, Chinese, Japanese in kanji, Turkish, Hebrew, and Arabic.

Keyboards for international models vary in layout, but, whatever the key arrangement, the character codes are mostly the same. Most screen and PostScript fonts for international Macs have all the common characters found in most Western European languages. A few companies produce fonts with the diacritical marks needed for some Eastern European languages as well.

Many programs are available in several languages. If you share files internationally, look for translated software in each language of interest. A Macintosh file, such as one from the international versions of Microsoft Multiplan, Chart, File, and Excel, "remembers" the language in which it was generated. If you use the German version of Microsoft Excel to open a spreadsheet created with the French version, for example, the program will ask you if you want to keep or change the currency used. Microsoft Chart also asks if you want to enter a currency exchange rate and rescale the values on graphs.

PRIVACY

Many people are much more careless with their disks than with their papers. If you have confidential information in a computer, keep it on a floppy, not a hard disk, and lock up the floppy when you aren't using it. If you must put a sensitive file on a hard disk, erase it when you have finished working with a utility program that not only changes the directory entry but magnetically wipes the tracks clean as well.

You can also use a program to encrypt a file so that a password is needed to read it. You must go through the decryption and encryption process every time you use the file, however, and if you forget your password, you've lost the file. If you are the only person who has access to that file, locking the disk it is on in a drawer is simpler and quicker than encryption, but if several people use the file, or if it must be sent somewhere, encryption is better.

One encryption program is Sentinel, version 1.0 (SuperMac, Mountain View, CA). The program is fast, but the icons of encrypted files are still visible in the disk window, so the files can be accidentally or deliberately tossed in the Trash. MacSafe, version 1.08 (Kent Marsh Limited, Houston, TX), is much slower than Sentinel, but it lets you hide the icons of encrypted files so that no one can tell they exist.

Encryption programs can also help with electronic mail. If only you and the recipient know the password, spies cannot decode your message even if it is on a public electronic bulletin board. Of course, encryption schemes can be broken, but they will deter most snoopers. The encrypted files created by MacSafe and Sentinel cannot be sent through a text-only electronic mail service, but they can be sent via MCI Mail with the communication program Desktop Express (Dow Jones, Princeton, NJ).

If your computer is connected to a network, disconnect it for privacy, especially if your network software makes every storage device available to the entire network. Password protection on network file servers is often weak; the files themselves are not usually encrypted, and a list of authorized passwords has to be stored somewhere so that the file server can verify the identity of users signing on to the network. If you must store a sensitive document on a network device, encrypt it first.

Intrusions to personal computers through telephone lines are rare because microcomputers are not usually left permanently connected to the phone line, and callers cannot generally gain access to a file directory or the operating system. If you do set up an unattended Mac for remote access, you should at least have a password. If you run into trouble, use a system that keeps an audit trail of activity and automatically disconnects on repeated efforts to gain access.

The best simple protection against intrusion by phone comes with call-back modems. Calls to such a modem are greeted with silence; the caller must then enter a touch-tone code number and hang up. The modem checks this code number against an internally stored list of numbers and dials the caller back, this time establishing a modem connection. Since the modem only dials numbers it already has in memory, the system is hard to break into. Call forwarding can be a weak link, however; consult data security experts if the stakes are high.

Most computers generate spurious electromagnetic radiation that a sophisticated snooper can interpret. Careful shielding can reduce or eliminate this radiation. (The military Tempest specification is a maximum spurious-radiation standard.) Again, you should consult with data security experts if you need such protection.

Also, watch what you throw into trash cans. Most printer ribbons create no security problem, but single-strike carbon ribbons, such as the IBM Selectric ribbon used by the Juki 6100 printer, can be read after use. Finally, if your hard disk needs service, think about its contents before sending it out; you may find it's cheaper to replace the disk than to leak its contents.

Serious fearmongers can go on and on, but we will stop here.

TROJAN HORSES AND VIRUSES

As their names suggest, computer Trojan horses and viruses are programs or sections of code that invade a host computer and damage or destroy data. Propagated by malevolent pranksters, these new threats to computers are as treacherous as their namesakes.

A Trojan horse can appear innocent enough; it might be a public-domain utility program on an electronic bulletin board or what looks like a legitimate shareware program whose code has been invisibly altered. When you start it, the program erases or damages your disk files instead of doing what you thought it would do; the program might alert you ("Ha! Gotcha!"), or it might not. A really insidious Trojan horse may do no damage the first few times you run it; only on the tenth or fiftieth running—after you might have copied it to other disks—does it start trashing your files.

Responsible managers of electronic bulletin boards (SysOps, or system operators) periodically check programs on their boards to rid them of Trojan horses. If a program arouses your suspicions, don't risk putting it on your hard disk; test it first on floppies. If you have a Mac SE or Mac II with a built-in hard disk, keep the hard disk off your Desktop by pressing Option-Command-Shift-Backspace (Delete) and restarting the machine from a system floppy disk.

A computer virus is even more dangerous than a Trojan horse because, much like its biological namesake, it can spread from computer to computer of its own accord. A computer virus is a short section of programming code that can attach itself to messages transmitted over computer communication networks spanning the continent. Once a host computer is "infected," the virus reproduces itself, and the infected messages spread throughout the network.

A virus may become part of a computer's operating system, which means it can spread every time you do something as routine as copying a file to another disk. Complex operating systems, particularly those with many built-in communication functions, such as UNIX, are more vulnerable to viruses because there are more potential places where they can enter. Simpler operating systems like the Mac's are less vulnerable but hardly immune.

HINTS FOR THE DISABLED

This section offers a few suggestions on how the Macintosh can be adapted for use by people with physical disabilities. For more ideas and information about products for disabled computer users, contact the organizations listed at the end of this section.

For the Hearing-impaired

If you are deaf or hard of hearing, you will generally have little difficulty using a Mac. You can set the speaker volume to maximum with the Control Panel desk accessory or use an auxiliary amplifier and speaker or earphone. Except for the few programs specifically designed to produce sound output, the speaker is used only to sound warnings, which are nearly always accompanied by a dialog box or an alert box on the screen. Alternatively, if you set the volume control to zero, the speaker will be silenced, and the system will alert you by flashing the menu bar briefly from white to black.

TDD/Baudot terminals (sometimes called TTY) are widely used by the hearing-impaired to communicate by typing. To use a Macintosh as a TDD terminal, you will need a communication program and a 110-baud modem that can handle 5-bit Baudot code. These items should be available from independent vendors.

A Macintosh cannot be practically used to decode closed-captioned television programs; the cost would be higher than the present decoders available for television sets.

For the Vision-impaired

The Macintosh can be adapted more readily than other microcomputers to work for people with certain forms of visual impairment.

For people whose visual acuity is limited, the computer equivalent of books set in large type is easy to achieve. One method is to connect a bigger screen to the Mac; this should be a screen with the same number of pixels as a standard Mac screen. (See the illustration on page 56). Another method is to enlarge the image itself. InLarge (Berkeley System Design, Berkeley, CA) lets you choose a magnification of 2 to 16 times. Once magnified, the image is larger than the screen; to see the entire image, you pan with the mouse. To help locate the insertion point, inLarge draws a cross hair across the screen centered on it.

LoVE (Low Vision Editor) from Rosetta Software (Newport Beach, CA) and Textalk (Assembly Corner, Maynard, MA) are two word-processing programs for the visually impaired that combine large type on the screen with speech synthesis.

Talker (Finally Software, Newport Beach, CA) is a speech synthesis word processor. The speech synthesis feature helps anyone who cannot read a video display. Because the speech synthesizer can read any standard ASCII text, the Macintosh could serve as a general-purpose reading machine.

If you are totally blind, you may find a Macintosh with its visual interface harder to use than a computer with a command-line interface that can be operated easily from the keyboard alone. If you live or work with people who prefer a Macintosh, however, you can still use it with a utility program to operate the computer without seeing the screen. MacPS (Neff Systems Group, San Jose, CA) automates Finder operations, such as starting programs and opening files. Tempo (Affinity Microsystems, Boulder, CO) is a keyboard macro program that can record a series of operations and assign the whole series to a single command key or key combination.

Brown Disc (Colorado Springs, CO) has announced a line of floppy disks that come with braille labels. According to the company, the labels for 3½-inch disks will fit through the narrow disk drive slots of a Macintosh.

Braille printers can be connected to a Macintosh in the same way that a daisy-wheel printer can, and braille keyboard overlays are available from Access Unlimited-SPEECH Enterprises (Houston, TX).

For the One-handed

Shift-key operations—those involving the Shift, Option, Command, and Control keys combined with another key on the keyboard or the mouse—usually require two hands. To do them with only one hand, try Apple's Easy Access, a keyboard-enhancing INIT program that lets you set the shift keys so that you can press key combinations in sequence rather than simultaneously. For example, to get italics in Microsoft Word, you can press Command, then Shift, and finally, the *i* key, instead of pressing Command-Shift-i all at once.

For Those with Motor-Control Limitations

If you have difficulty with fine motor control, you might have trouble using the mouse effectively. Sometimes you can remedy the problem by modifying the mouse driver to change the mouse's "gearing." A modified driver could also be designed to register only smooth movements and ignore sudden or jerky motions.

Some users may find that moving the mouse button to a different part of the mouse, or out of the mouse altogether, makes it easier to operate. In some cases, a trackball, joystick, digitizing tablet, or head-mounted pointing device may help. (See Chapter 5.) The Easy Access program also gives you complete control of mouse movements from the keyboard.

If you are operating the keyboard with only one finger, a headstick, or a mouthstick, a keyboard guide with openings for each key may help to prevent you from hitting keys accidentally. Prentke Romich Company (Wooster, OH) and Adaptive Technology (Brooklyn Center, MN) make such guides.

Sharing Good Ideas

If you have found a useful way for a disabled person to use a Macintosh or any other computer, publicize it. Write an article for a magazine or a letter to the editor. Put suggestions on electronic bulletin boards and pass them on to organizations for the disabled.

The following are some information sources about computers and the disabled:

- SpecialNet is a national communications system with more than 40 different electronic bulletin boards that discuss special education issues; it is widely used by special education and rehabilitation personnel. Subscriptions are $200/year. For information, contact National Systems Management, Inc., 2021 K Street NW, Suite 315, Washington, DC 20006, telephone: 202-296-1800.
- Apple is supporting the National Special Education Alliance to help disseminate microcomputer technology to aid disabled individuals. To locate the center nearest you, contact the Office of Special Education, Apple Computer, 20525 Mariani Avenue, Mail Stop 36M, Cupertino, CA 95014, telephone: 408-973-3854. Apple itself maintains Special Education Solutions, a database of products for the disabled; access to the database is available through Apple dealers.
- The Center for Special Education Technology is an information clearinghouse that operates within the Council for Exceptional Children, 1920 Association Drive, Reston, VA 22091, telephone: 703-620-3660.
- The Trace Center publishes the International Software/Hardware Registry, a resource guide for disabled students. Contact the Trace Research and Development Center, University of Wisconsin, 314 Waisman Center, 1500 Highland Avenue, Madison, WI 53705.
- The newsletter *Closing the Gap* is another good information source; it can be reached at Post Office Box 68, Henderson, MN 56044.
- For information about using the Macintosh to teach dyslexic children, write to Richard Wanderman, Director, Computer-Aided Writing Project, The Forman School, Litchfield, CT 06759.

THE MACINTOSH VERSUS IBM PCs

When IBM introduced its PC in 1981, Apple was the foremost maker of personal computers. By the end of 1983, IBM had taken the lead, riding on a flood of software and attracting dozens of imitation PCs from other manufacturers. Probably no one, not even IBM, anticipated such success. Today, the sheer number of models complicates any comparison of Apple's offerings with IBM's line. In this section, the term *PC* refers to IBM microcomputers as well as clones, or copies of IBM models.

Apple and IBM developed their microcomputers from almost opposite design philosophies. The Macintosh family is built around an elegant concept, the visual interface, that is carried through all its software. The best Macintosh programs—written by Apple as well as by independent firms—are the best for any computer.

The IBM PC, on the other hand, follows the long-established tradition of general-purpose computers: It is a hardware base upon which software developers must build an interface as well as application programs. Nearly all the useful software for the IBM PC series has come from independent companies, and each of these went off in its own direction, creating chaos. It's hard to transfer even simple text between programs for IBM PCs; transferring graphics is extremely difficult. The microcomputer software that has been developed by IBM itself is mediocre. Indeed, in its entire history IBM has never produced great software for any computer.

With their 1987 models, Apple and IBM have moved toward each other in design. Apple now produces more flexible hardware, and IBM has announced that it intends to adopt a software interface similar to the Macintosh's. IBM's plans require nothing less than rewriting all existing PC programs, a process that will take years. Until such rewriting is done, IBM PC users must make do with old-fashioned software.

In terms of hardware, a Mac SE has raw processing power comparable to an IBM PC/AT or PS/2 (Personal System) series 50, micros based on the Intel 80286 CPU. A Mac II is comparable to a COMPAQ 386 or IBM PS/2 Model 80, microcomputers based on the Intel 80386 CPU.

The IBM PC world has some important hardware advantages over the Macintosh, including IBM-made keyboards, which are far superior to any keyboard available for the Mac. In addition, effective portable PC clones are available from several manufacturers, as are low-cost desktop PC clones.

PC clones are easy to make, and their manufacturers do not require permission from IBM; stiff competition has driven many companies to both reduce costs and design distinctive features into their clones. Apple, in contrast, has no direct competition for the Macintosh, so it can build a few models that it sells for relatively high prices, giving the prospective buyer few choices. A Macintosh clone is possible, but much more difficult to build than a PC clone, unless Apple becomes willing to license its circuits.

The real differences between PCs and the Macintosh stem more from software than from hardware. The PC has more software titles than the Macintosh, particularly in specialized business domains; in certain areas, such as graphics and desktop publishing, the Mac has more. Nearly all PC software is poorly designed and harder to learn and use than software for the Mac. PC clones may be cheap, but the cost savings could be offset by the amount of time it takes to learn how to use one—much longer than training time on a Macintosh.

To make matters worse, the operating systems and the software for the PC are splitting into three incompatible variants: MS-DOS and software that runs on the Intel 8088 and 8086 CPUs, a version of OS/2 (a new PC operating system) and application software that runs on the 80286 CPU, and another version of OS/2 and software that runs on the 80386 CPU.

Software for these different CPUs is only upwardly compatible. Software for the 8088/8086 will run on the 80286 and 80386, but not the reverse. Although 8088/8086 software will run a little faster on the faster chips, its performance is still constrained by the same memory limitations as on 8086/8088 machines; MS-DOS simply cannot handle more than 640 KB of memory. Software for the 80286 will run on the 80386, but software that will run on the 80386 has been tailor-made to take advantage of this chip and will not run on the 80286.

These incompatibilities arise from uncoordinated design decisions that Intel made for its CPU chips. The Motorola 68000, 68020, and 68030 CPU chips used in the Macintosh family are much more compatible, so Mac software is spared drastic upheaval when new hardware appears.

Alternatives to Apple and IBM do exist. If Macintosh prices are too high for you, but you want a microcomputer now that runs modern software, Atari's ST and Commodore's Amiga are possible choices. Although ST and Amiga software is far less developed than the Mac's, the hardware is cheap.

19: Advanced Communications

This chapter examines the basis of computer communications and then looks at some specific topics, such as using a Macintosh as a terminal to a large computer and facsimile.

THE BEGINNINGS OF CHAOS

Communication between Macintoshes may not be difficult, but problems begin when you try to use a Mac to communicate with other kinds of computers. The problems have nothing to do with Apple or the Mac specifically; the headaches belong to the whole computer industry. Occasionally you can get a transfer between computers to work quickly; more often it takes hours, sometimes even days, to establish a connection for the first time.

The problems lie in the profusion of protocols (specific forms of the communication signal) and hardware standards, and nonstandards, that litter the computer industry. Adding to the confusion, the words that describe communications are often used ambiguously.

The ISO Protocol Layers

Considering how complex communication is in the world in general, the chaos in the computer world may not seem so surprising. In speech, for example, we follow protocols about who should speak first, forms of address (Mister, Ms., Sir, Madam), ways to keep others from talking ("...uh...," "You know,..."), and so on. For most of us, these protocols are ingrained as part of our social training. Keeping the many levels of computer protocols straight, on the other hand, requires some concentration.

The International Standards Organization (ISO) has developed a model with seven protocol layers to describe communication between different computers. The following explanations of these layers include computer examples—which will be explained later—and analogous concepts from other forms of communication, principally the postal system. The list starts at the lowest hardware layers and works up to the highest software layers. Each successive layer builds on the ones preceding it.

The physical layer covers the mechanical, electrical, and functional arrangements necessary for a physical connection. Examples: Computer cables, telephone lines. Analogies: Trucks, railroads, postal delivery personnel.

The data-link layer is used to package and send a unit of information from one node to another. A node is any device—a computer or printer, for example—that is able to send or receive information, or both. This protocol includes flow control—who sends data and when—and some forms of error control. Examples: An asynchronous serial connection, a 1200-bps modem. Analogy: An envelope.

The network layer determines how information from the sender is routed to the correct receiver. This step is usually missing if only two devices are involved. Example: A token-passing network. Analogies: The address and the mail-sorting process.

The transport layer includes the steps taken to ensure high-quality network service, including confirmation that the information has reached its destination and has been read without error. This layer is not always used; many computers simply send out information without knowing where it goes. Examples: The error-detection and correction features of many communication programs. Analogy: The sender requesting a delivery receipt from the post office or acknowledgment from the receiver.

The session layer contains the procedures by which two communicating computer devices coordinate action. Example: The sending computer requesting that the receiving computer open a file, accept information, and close the file; the person at the receiving computer need not intervene. Analogy: The sender asking the receiver to perform an action they both understand.

The presentation layer protocol details all the formatting and code conversion necessary to make information from the sending computer intelligible to the receiving computer. Example: Supplying graphics in a form that the receiving computer can display; the SYLK or Lotus WKS data file formats. Analogy: A lawyer sending a letter to a client that explains a court decision in plain English.

The application layer protocols ensure that the information is sent in a form that can be used directly by an application program running on the receiving computer. Example: Moving a spreadsheet file from an IBM PC to a Macintosh. Analogy: A lawyer sending a legal analysis of a court decision to another lawyer.

There is no shortage of computer protocols at each of the seven layers; moreover, many protocols cut across several layers. Indeed, there are too many protocols, and most of them are incompatible with one another, so things are not nearly as tidy as the above list implies.

A given computer communication link requires agreement between sender and receiver at some, but not necessarily all, protocol layers. For your computer and another computer to communicate at all, you and the other user must agree on the physical and data-link protocols; and for the computers to communicate smoothly, you must agree on the transport and session protocols. For fully integrated operations—where both sender and receiver can interact directly with the transmitted information—you must agree on the presentation or application protocols.

THE PHYSICAL LAYER: CONNECTING THE HARDWARE

A communication channel can be full-duplex (supporting simultaneous, two-way communication), half-duplex (communication is one way at a time, but the direction can be changed), or simplex (one way—and one direction—only, such as a radio broadcast). Simplex is rarely used for data communication; whether a link is full-duplex or half-duplex depends on the hardware and protocols.

Direct connections tie two computers together, via their serial ports, with a simple cable. The cable, which replaces a pair of modems, is sometimes called a null-modem because no modem is needed. (See the Appendix for wiring diagrams of Mac-to-Mac and Mac-to-other-computer null-modems.)

Connection via a local area network (LAN) electrically links many computers and accessories, such as printers and high-capacity hard disk drives. Unfortunately, two devices that can be physically attached to the same network may not be able to talk to each other unless the upper-layer software protocols permit. Better networks accept a variety of computers; restrictive ones accept only a single brand or model. AppleTalk is a simple, low-cost network; many other types are available. Some can move data much faster than AppleTalk, but they are also more expensive.

Connection via telephone lines requires a modem attached to each communicating computer. Both computers must be ready at the same time; one then calls the other. Depending on the software and hardware, an operator may not need to be at the answering computer or even at the sending computer; everything can be run by a timer under software control.

The telephone line may be:

- Public switched telephone network (PSTN), the same one used for routine phone conversations. The quality of these voice-grade, dial-up switched lines varies greatly and appears to be getting worse since the deregulation of telecommunication. Under good conditions, PSTN's channel capacity is about 20 kilobits per second (Kbps).
- Leased, or dedicated, lines with higher bandwidth and less noise than PSTN. Because these lines are tested for performance and sometimes have

permanent electrical compensation added, they are also called conditioned lines. Leased lines come in a variety of channel capacities.

- Integrated Services Digital Network (ISDN), which will eventually replace ordinary dial-up and dedicated phone lines. Businesses are now starting to use ISDN; conversion of residential phone lines is expected to begin early in the twenty-first century. ISDN encompasses many different services, supporting everything from computers to voice and, ultimately, television. Each channel type is assigned letter codes by the International Telegraph and Telephone Consultative Committee (CCITT).

  The most widespread ISDN service runs on a single telephone line; it is called 2B+D, meaning that it has two 64-Kbps B channels that can carry digitized voice or data and one 16-Kbps D channel carrying digital control signals. With additional phone lines, multiple B channels can be combined for higher capacity; the basic configuration for a PBX in North America, for example, is 23B+D with a total data rate of 1.544 Mbps (megabits per second); in Europe it is 30B+D (2.048 Mbps).

  For the average computer user, the 64-Kbps B channel will become the new standard, offering much faster transmission speed than any modem on a PSTN line. Because ISDN is fully digital, no modem is necessary; the computer connects to the phone line through a simple interface circuit. (Even so, a 64-Kbps B channel is not fast enough to support many forms of data transmission, such as high-fidelity sound or television. The forthcoming standards for broadband ISDN incorporate 150-Mbps channels as building blocks. Five 150-Mbps channels, for example, could support digital high-definition television.)

Radio transmission can also furnish a computer link. Wireless modems typically operate at 72 MHz, a frequency set aside for business communication, with a maximum range of about a mile under good conditions; Electronic Systems Technology (Kennewick, WA) supplies such modems. Cellular radio works but requires a robust error-detection and correction scheme, such as the one developed by Spectrum Cellular (Dallas, TX). Air-to-ground and ship-to-shore radio links are too unreliable for any but the most determined users with the best error-correction protocols. Data transmission via amateur radio is permitted on some frequencies by the FCC; reliability varies tremendously with location, time, and frequency.

Infrared links operate by means of a modulated infrared signal—somewhat like a television's remote control—to send data within a room.

Connection via a third computer may work when connection between two computers proves too elusive or schedules don't match. The first micro communicates information to a third computer; the second retrieves the information. Most electronic mail services follow this pattern. Because an intermediate computer is used for storage, the two micros need not be available at the same time. The third computer is usually a large one, often able to handle several kinds of protocols or at least one well-defined protocol. Getting a micro to communicate with a large computer can be easier than getting two micros to talk directly with each other. This method is usually restricted to text, although you can code any file to a format that looks like text as far as the third computer is concerned.

THE DATA-LINK LAYER: CODING THE SIGNAL

The form the signal takes in order to pass through the physical link—and whether communication is asynchronous or synchronous—depends on both hardware and software. In asynchronous ASCII communication, which is the most common kind, each character is sent separately, preceded by a start bit and followed by one or two stop bits and possibly an error-detection bit. Because the receiver can always tell from the start and stop bits when a character starts and when it is complete, the sending and receiving computers do not have to be synchronized to the same clock—in other words, they are asynchronous, with variable timing between characters. They must, however, work at the same speed.

Putting in all the start and stop bits takes time, however. Transmission rate improves considerably if the sending computer transmits not merely a character at a time but a block of characters (typically 256) without start and stop bits. Characters are delimited by timing instead. To transmit a block, the sending computer puts synchronizing information at the beginning of the block. The receiving computer uses this information to measure off each character in the block. The two computers then operate synchronously, with fixed timing between characters and variable timing between blocks.

If you are typing while using a synchronous communication link, and you are a slow typist, a block may contain only one character; the rest of the block simply marks time. In some systems, block length is variable; this is a valuable feature for interactive communications (where short blocks speed things up) and for coping with noisy transmission lines (short blocks help with intermittent noise; normally, long blocks are more efficient because the overhead is relatively less).

Synchronous communication is used mainly with mainframe computers and computer networks, such as AppleTalk. Microcomputers generally communicate asynchronously because it is easier, and the performance penalties are not too serious for small quantities of data. We will therefore concentrate on asynchronous communication to continue our discussion of the protocol layers; we will return to synchronous communication later.

Asynchronous Communication

First, you need communication software at each end. If possible, get software for each computer from the same company. Although the programs themselves may differ so that they can run on different computers, they are more likely than programs from disparate companies to be compatible at several levels—sometimes as far as the session layer protocol—provided that both programs are set up with the same parameters.

If you cannot get such software, then you must configure two different communication programs to use the same data-link protocol. (Some programs follow their own protocols at the transport layer to send and receive and will therefore work only if the other machine is equipped for the same protocols.)

The protocols at several levels may restrict the kinds of information you can send. Between different types of computers, the transmitted information is most often ASCII text. Although, strictly speaking, ASCII code includes all possible 7-bit characters, in communication parlance, an "ASCII file" normally contains printing characters—letters, numbers, and punctuation—plus tab, carriage return, linefeed, and form feed. The remaining characters, including most control characters, are normally nonprinting. Many communication programs transmit some control characters; Control-g is useful because it will ring a bell on most terminals. (Most software interprets a Command-g on the keyboard as Control-g if the keyboard lacks a Control key.) Other control characters, such as the end-of-file marker, are usually removed before transmission.

Even if all you want to do is send an ASCII file, both machines must agree at the data-link level. The hardware determines some parameters, the software determines others; still other parameters can be set by either. The most important parameter is speed.

Speed

Baud is a term from nineteenth-century telegraphy for how fast information is sent. Strictly speaking, one baud is one signal event or modulation change per second (a "symbol" in communication theory). At 300 baud, one baud—one signal event—corresponds to one bit per second; at higher speeds, one signal event codes two, four, or more bits. A "1200-baud" modem actually sends 600 events per second; each event codes two bits.

Although *baud* is commonly used interchangeably with the term *bits per second,* or bps, the strict definitions are followed here. Table 19-1 gives, among other things, the data rate and the true baud rate for all common modem types. Depending on the exact data-link layer protocol, 300-bps transmission carries 30 to 33 characters per second; 1200 bps carries 120 to 132 characters per second; 2400 bps carries 240 characters per second. By coincidence, the number of English words per minute is about the same as the number of bits per second.

Although the bit rate can be continuously varied in some cases, the industry has settled on several standard speeds and hardware protocols, largely set by either Bell Telephone Laboratories (North America) or CCITT (first in Europe and Japan, now worldwide). If you use modems to communicate, they must operate at one of the common speeds and compatible hardware protocols. For computers directly wired together, you can select the highest speed permitted by your equipment and the length of the connecting wire; the longer the wire, the lower the speed. If in doubt, start at a low speed and increase it until you encounter errors; then back off one step. (If you start at a high speed and get errors, you cannot always tell whether they are caused by excessive speed or by something else.)

Low-speed modems

As faster and faster modems appear on the market, the definition of *slow* shifts upward; we consider the following modem speeds low:

- 75 to 110 bps: used by teletypes, Telex machines, other systems supporting older printers, and some radio applications. Also used with Baudot code instead of ASCII for TDD communications by the hearing-impaired; you'll need a translation program and a 110-bps modem to connect to TDD devices. The 5-bit Baudot code has uppercase characters only.
- 300 bps: the traditional home-computer speed. At 25 to 30 characters per second, these modems are slow, taking a minute to fill a Mac screen with characters. Most people can read text comfortably at this speed (up to 300 words per minute). Almost obsolete, 300-bps modems are giving way to 1200-bps modems. The standard protocols for 0-bps to 300-bps communication are Bell 103 in North America and CCITT V.21 in Europe and Japan; the two are not compatible.
- 1200 bps: most common in business and increasingly common with home computers. Transmission at 1200 bps is too fast for most people to read—a Mac screen fills in 15 seconds—but it can be scanned. If you're using a commercial database service that charges for each minute of use, save everything on disk and read it later; it's much cheaper.

  The standard 1200-bps protocols are Bell 212A in North America and CCITT V.22 in Europe and Japan; the two are not compatible. The Racal Vadic 3400 protocol is still used occasionally, particularly by computing centers; most modems that use the Racal Vadic form can also use Bell 212A. These three protocols are full-duplex and work on ordinary phone lines. Avoid products that use Bell 202, a lower-cost, half-duplex protocol.
- 2400 bps: the fastest common speed. The new worldwide standard for full-duplex communications over dial-up phone lines is the CCITT V.22 bis protocol. It is adaptive, adjusting both sending and receiving characteristics to phone-line conditions when communication begins. The V.22 bis protocol is complex, requiring clearer and more noise-free telephone lines than 1200-bps communication does; many long-distance connections aren't good enough.

  V.22 bis modems are supposed to fall back to the 1200-bps V.22 protocol if the line is too noisy, but some units don't, or they fall back instead to the incompatible 212A protocol. If you set out to connect at 1200 bps with a V.22 bis modem, you must determine whether it is set for 212A or V.22. When communicating with the same model V.22 bis modem, you should not have these concerns—if you are communicating with a modem on the same continent. But because of differences between European and American telephone practices, several potential problems can crop up when you try to communicate internationally with modems that nominally follow V.22 bis. For example, the European answer tone (2100 Hz) differs from the American one (2225 Hz); calls from America to Europe generally work, but the reverse isn't always true.

These problems arise from the way the standards were set. Bell Telephone actually built and sold a model 212A modem at a time when it controlled the American telecommunications network. Companies could buy it and check their own units against it. The CCITT, on the other hand, is a voluntary body that can only make recommendations. No V.22 bis modem standard exists, only a pile of paper that manufacturers can and do interpret in their own ways.

The CCITT V.26 ter protocol also runs at 2400 bps. It is technically superior to V.22 bis and can cope with noisier telephone lines. It is also much more expensive, using echo canceling rather than frequency-division multiplexing for full-duplex operation. V.26 ter units are also more compatible with V.32 4800-bps and 9600-bps modems.

Medium-speed modems

Modems that run between 4800 and 19,200 bps are generally used for synchronous communications, but they can be used asynchronously as well. The older V.29 modems are half-duplex on dial-up lines; the newer CCITT V.32 protocol defines modems that are full-duplex with echo canceling on dial-up lines. V.32 modems are incompatible with V.29 modems.

Because they are so expensive, full-featured V.32 modems have not been sold for microcomputers, but new integrated circuits should reduce costs significantly in 1988. In the interim, many manufacturers have introduced 9600-bps modems that use either V.29 or a simplified half-duplex subset of V.32 protocols. These various stopgap modems are mostly incompatible with one another.

Telebit (Cupertino, CA) makes a modem that operates at up to 14,400 bps on ordinary dial-up phone lines; ingenious technology enables it to send as many as 512 carriers simultaneously, with inherent error correction. Like some 9600-bps modems, this Telebit modem can also operate at 2400, 1200, and 300 bps.

Table 19-1 gives characteristics of all standard modem configurations up to 9600 bps. Systems that are half-duplex on a two-wire telephone line can often be turned into full-duplex with four wires or two telephone lines. Leased and dial-up lines are to some extent interchangeable; the table shows which is most common for each protocol. Besides the modem types described here, there are many other proprietary protocols used by single manufacturers and suitable only for closed communication within an organization.

In any country outside of North America, Europe, and Japan, which protocols are used depends mostly on who are the country's main economic partners. Many countries still officially ban any connections to the phone except by the central telephone company. Telephone-line quality varies widely; some phone systems are unusable by modems at any speed.

TABLE 19-1. MODEM PROTOCOLS

Protocol	*Data rate (bits/sec)*	*True baud rate (signal events/sec)*	*Half-/full-duplex*	*Duplex type*	*Async/sync*	*Dial-up/leased lines*	*Area*
Bell 103	300	300	F	FDM[1]	A	D	North America
CCITT V.21	200–300	200–300	F	FDM	A	D	Europe/Japan
Bell 202	1200	300	H	—	A	D	North America
Bell 212	1200	600	F	FDM	A/S	D	North America
CCITT V.22	1200	600	F	FDM	A/S	D	Europe/Japan
Bell 201	2400	1200	H	—	S	D/L	North America
CCITT V.22 bis	2400	600	F	FDM	A/S	D	All
CCITT V.26 ter	2400	1200	F	ECT[2]	A/S	D	All
Bell 208	4800	1600	H/F (2/4 wire)	—	S	D/L	North America
CCITT V.32	4800–9600	2400	F	ECT	A/S	D	All
Bell 209	9600	2400	H/F (2/4 wire)	—	S	L	North America
CCITT V.29	9600	2400	H/F (2/4 wire)	—	S	D/L	All

[1]FDM: Frequency-division multiplexing
[2]ECT: Echo-canceling technique (see text for explanation)

High-speed modems

Traditionally, modems faster than 19,200 bps have been used for specialized applications such as linking two mainframe computers. Some models can operate at a megabit per second over short distances. Special digital communication links may use microwave or television circuits. Few computer users work directly with this class of modem.

Noise on the Phone Line

Noise limits the capacity of any communication channel. The standard predivestiture AT&T phone line, measured under modem communication conditions, had a 24-db signal-to-noise ratio (S/N) on long-distance lines and a 27-db S/N for local calls. Table 19-2 lists the S/N required for different modem protocols. (The figures were compiled by Ken Krechmer of Action Consulting, Palo Alto, CA; see his article in *Data Communications,* April 1985.)

TABLE 19-2. SIGNAL-TO-NOISE RATIOS

Speed	Protocol	Best-case S/N (db)	Typical S/N (db)	Safety margin with typical S/N (db)
300 bps	103	4	6	16
1200 bps	212A/V.22	7	14	10
2400 bps	V.22 bis	14	21	3
2400 bps	V.26 ter	10	17	7
4800 bps	V.32	16	21	3
9600 bps	V.32	19	Not available	Not available

In the table, best-case S/N gives the theoretical performance of a protocol; typical S/N is the measured performance for commercial products. The safety margin is calculated for a predivestiture long-distance line. The best industrial-grade modems perform distinctly better on noisy lines than mass-produced consumer units, but on good telephone lines there is no functional difference. Unlike analog signals (such as hi-fi sound), digital communications are not affected by noise as long as the bits are demodulated accurately.

The real situation may be worse than the table shows, however. Simple signal-to-noise ratios do not account for many problems, such as pulse noise. Furthermore, consumer modem quality varies widely; some cheaper designs have much worse S/N figures than the average and can work only on local telephone lines at 1200 bps. Some modems contain design errors that degrade performance in communications with a different modem model.

The signal-to-noise ratios given here measure only one-way links. Working modems on full-duplex links need an additional 3 db, effectively wiping out the safety margin of V.22 bis. The competition to offer telephone services at lower rates has already degraded signal quality on many long-distance lines, and the situation could get worse. With 2400-bps and faster modems, connect via a local call whenever possible; call a nearby gateway for a dial-up service rather than try to communicate directly over long-distance lines.

You have only two choices with noisy phone lines: Switch to a slower speed or adopt error correction (described in the section beginning on page 302).

Although measuring modem performance is complicated, a few magazines have done so for typical telephone circuits; see, for example, *PC Magazine,* May 12, 1987.

Other Parameters

After matching speed, both sides must also agree on other data-link layer parameters:

Character width or data bits defines the number of bits in a character: either seven or eight. If you are communicating with another Macintosh, eight bits can carry the entire character set; seven bits cannot. For ordinary text, the number doesn't matter as long as the communicating parties agree.

Stop bits—one, one and a half, or two—signal the end of a character. This setting frequently doesn't matter either; after counting off the data bits, a computer can ignore everything until the next start bit.

Parity—even, odd, or none—is a simple form of error detection. A single parity bit lets the systems determine if there is a one-bit error in the received signal. For even parity, the computer adds up all the bits in a character; if the sum is even, the parity bit is 0; if the sum is odd, the parity bit is 1. Either way, the final sum is always even. If a one-bit (or three-bit or other odd-numbered) error occurs, the sum will no longer be even. A two-bit error will not be detected because the sum remains even. Odd parity works the same way except that the sum is always odd. Some systems specify Mark or Space, which are variations on no parity.

What the receiving system does on encountering a parity error depends on the software. The parity bit is simply an error flag letting programs signal for retransmission—if the sending computer is set up to recognize the signal.

Duplex determines whether information travels one way at a time or both ways at once. At 300, 1200, and 2400 bps, most phone links are full-duplex—information goes both ways simultaneously. On a full-duplex link, the communicating computers use different frequency bands (frequency-division multiplexing). So that both don't try to use the same band, one is set to the *originate* band, the other to the *answer* band; which computer is set to which band doesn't matter as long as both sides agree. By convention, all modems and software set the calling party to the originate band.

At fast communication speeds, the telephone line doesn't have enough bandwidth for reliable transmission with two distinct frequency bands, so higher-speed modems use echo canceling, a technique that allows both modems to use the entire frequency band. Each modem senses the effects of its own transmission and cancels them, leaving the incoming signal free of interference. Many long-distance phone lines have their own echo canceling, which may interfere with the echo canceling built into a modem; check with your long-distance supplier if you think you have a problem. International Telephone and Telegraph has proposed that echo-canceling modems send a special switching tone that tells the long-distance service to turn off its own echo canceling.

With a half-duplex link, information travels only one way at a time. The line can be reversed with signaling codes, but frequent reversing slows transfer rate considerably. If an application calls mainly for one-way transfers, however, half-duplex presents little speed penalty over full-duplex.

Character echo lets you know that your characters are getting through when you are communicating with a remote computer. With full-duplex and character echo turned on, each character you send out is returned (echoed) by the distant computer. Echoing is only practical with a full-duplex protocol (it takes too long in half-duplex), but not all computers operating in full-duplex will echo. (Character echo is a completely different issue from echo canceling.)

If your echoed characters are occasionally garbled, your phone line is probably noisy. If the remote machine cannot echo, your computer must generate the characters on your screen (local echo). If all your characters appear doubled, turn off local echo; if you don't see anything on the screen, turn on local echo. (Some software and publications confuse duplex with echo.)

For the four data-link layer parameters discussed below, communications will be more efficient if sender and receiver agree, but if messages are short, or if the receiver is willing to rework the received information later, complete agreement may not be absolutely necessary.

Handshake is a way in which one computer can tell the other to pause. Sometimes the sending machine transmits faster than the receiving machine can receive. To set up a handshake, the receiving machine first sends a signal (usually DC3 or XOFF, Control/Command-s) to the sending machine, which then stops until it receives another signal (usually DC1 or XON, Control/Command-q). Handshakes help when the receiving computer must write the transmitted information on disk. Unfortunately, many senders do not recognize handshakes. (On short communication lines, such as one between computer and printer, a separate electrical circuit in the connecting cable may carry the handshake signal; this is called a hardware handshake.)

Line ends signal the ends of each transmitted line of text. Standard Macintosh programs store text files with a carriage return at the end of a line. The carriage return signifies a move back to the left margin and a one-line advance on paper and on the screen. Some other computers distinguish between a carriage return and a separate linefeed and require both. Good Macintosh communication software should add linefeeds to carriage returns if necessary and strip incoming linefeeds when accompanied by a carriage return. If these steps aren't taken, you may see double linefeeds in some cases, and the other party may see all your lines overwriting themselves.

Filters remove certain characters sent from some systems; without filters you may see spurious characters on your screen.

Pauses are required at the end of each line by some receiving computers to allow them to process the line.

THE NETWORK LAYER: ADDRESSING THE MESSAGE

This protocol layer generally affects only local area networks, specifying the addressing scheme so that a message goes only to the intended recipient. The layer is built into network hardware and software; it is part of AppleTalk software, for example.

THE TRANSPORT LAYER: ERROR-CORRECTION AND VERIFICATION PROTOCOLS

Protocols that transmit standard ASCII files in a continuous stream cannot send every possible combination of bits because some codes must be reserved for marking the end of a file or signaling for the handshake. Transport layer protocols get around this difficulty by sending information in blocks of a fixed size called protocol data units. Thus, because sending and receiving computers both know exactly how many bits will be in a block, all bit combinations are permissible within the block. The longer the file, the more blocks it uses.

Files that can contain any combination of bits are called binary files because they are simply a sequence of binary digits. Programs that run directly on a computer are binary files, as are many types of data files.

For error detection and correction, the sending machine calculates an error-detection code from the data in the block and adds the code to the end of the block. The receiving machine independently calculates the error code for each block and compares it with the code sent. If the codes don't match, the receiver requests the block again. Because of this additional information, the transfer rate is slower than sending straight ASCII without error correction. In some cases, error detection is done at the data-link layer, and correction is done at the transport layer. Transport layer protocols with error correction are sometimes called protocol transfers or verification protocols.

The computer industry has no standard transport layer protocol; several dozen are in use, many for specific brands of hardware or software. Many work in a primitive way: Before receiving information, the receiving machine must open a file and let the sending machine know that it is ready to receive; the sender then transmits the information, and the receiver closes the file.

The transport layer protocol best known among computer hobbyists has been Xmodem, used by many microcomputer communication programs for many different computer models. Xmodem is effective (although it does not catch all errors) but works only at the transport layer. Xmodem is not particularly efficient, for it works in a half-duplex fashion. After sending each block, the sender waits for acknowledgment; on satellite transmissions, this takes a long time—80 percent of transmission time at 2400 bps. In faster protocols, the sender transmits continuously, getting acknowledgments from the receiver while transmitting. When errors are detected, the sender can send a block again without pause. Xmodem's fixed block size makes it unsuitable for interactive communications when you are sending only a few characters at a time—choosing items from a menu, for example.

Newer protocols, such as the Microcom Networking Protocol (MNP) class 4 and Link Access Procedure-Balanced (LAP-B), which is part of a protocol called X.25, have more robust error checking than Xmodem. Both MNP and LAP-B can work synchronously, saving time by eliminating start and stop bits (provided the modem can cope), and both have variable block size. The two protocols work most efficiently when they are built into the modem hardware. MNP and LAP-B are incompatible: An MNP modem connected to an LAP-B modem can communicate, but without error correction. A few manufacturers have announced plans to build both protocols into their modems.

MNP automatically checks to see if the receiving modem understands MNP and if the receiving computer has enough space on its disk to store the file to be transmitted; whatever the case, MNP can take the appropriate action. MNP class 5 includes data compression. (See page 305.)

LAP-B is an international standard, although fewer modem manufacturers use it than MNP. It is similar to error-correction protocols that are part of ISDN standards for both B and D channels.

The Kermit error-correction protocol developed at Columbia University also works at the transport layer. Because it was designed for universal use on any type of computer, it is inherently inefficient. Kermit is available for the Macintosh; it is used mainly to connect systems in universities.

When a Mac uses a transport layer protocol, the communication program usually sends only the data fork of a file. If you want to send a complete Macintosh file, use something like MacBinary, which is a protocol that overlaps the session layer and specifies a way to send and store a complete Macintosh file on other computers. Such files can be used for electronic mail and general storage.

THE SESSION LAYER: COORDINATING ACTION

Two communicating computers using compatible session layer protocols can automatically act—to open or close files, for example—on a command from one or the other. The transmission essentially moves not only the contents of a file, but the file folder itself, labeled and ready for use. Most local area network protocols work at this level.

In transmissions between Macintoshes, the session layer protocol sends a file's type, icon, and other attributes along with its contents. If the communication involves several computer types, session layer protocols can keep track of which computers understand which file contents. A Macintosh might receive a file with an icon, while an IBM PC might get only a directory name.

The first choice for a session layer protocol for Mac-to-Mac communication is MacBinary. It transfers a complete file, including headers and both data and resource forks. MacBinary is widely supported by Macintosh communication programs and has largely replaced Macintosh Xmodem, an earlier session layer protocol that is not directly compatible with the traditional Xmodem on other computers.

For communications between a Macintosh and an IBM PC, MacLink Plus (DataViz, Norwalk, CT) operates at the session layer; you can see and change IBM disk directories and subdirectories from the Macintosh. MacLink also performs some file format conversions. (See Chapter 20.)

THE PRESENTATION AND APPLICATION LAYERS

Presentation layer protocols enable the receiving computer to make use of information transmitted in a standard format, although the incoming information must still be converted into a format that an application program can work with. Presentation layer protocols range across many applications.

Graphics protocols describe a screen image. In the Macintosh world, QuickDraw is the most common format. Unfortunately, no industrywide screen graphics protocols exist. Text protocols mostly depend on file format standards. MacWrite files can be read by many other Macintosh programs and a few IBM PC programs. Lotus's WKS format for spreadsheet data can be read by many programs, including Microsoft Excel and Modern Jazz.

When there is agreement at the application layer, files can be moved between, and used directly by, both sending and receiving computers. Until there are more standards, this level of integration is achieved mainly when similar computers are communicating.

For specific information on moving files from another computer to the Macintosh, see Chapter 20.

COPING WITH DATA (OVER)FLOW

The Macintosh and other graphics-driven microcomputers create and manipulate far more information than non-graphics-based microcomputers. Files containing not only ASCII text, but also typefaces, icons, graphs, and drawings, gobble up time in transmission and run up the phone bill. If you frequently send large quantities of information, you should consider ways to increase your effective transmission rate, for example:

- Faster modems. Since higher-speed modems may not operate reliably on noisy dial-up lines, you may want to switch to leased lines. If so, the telephone line charges go up to about $1.50 per mile per month.
- Data compression and data concentrators. An English-language text file contains much redundant information. A trivial example is the letter *q*, which is always followed by a *u*. If sender and receiver agree, the *u* can be deleted from the transmission, saving time; the receiver inserts a *u* after every received *q*.

  The widely used Huffman coding algorithm compresses English text to about half its length; a complementary algorithm restores the original text. Compression ratios depend on the information; the more redundant the information, the tighter the compression can be. The most powerful compression techniques work with a variety of data, not only with English text. Compression speeds up data transmission simply by reducing the amount of information that needs to be sent.

You can compress files manually by running them through a compression program such as PackIt (shareware from Harry Chesley, 1850 Union Street #360, San Francisco, CA 94123) or the more sophisticated StuffIt (shareware from Raymond Lau, 100-04 70th Avenue, Forest Hills, NY 11375). Modems with MNP class 5 have compression built in, along with error correction. Some companies, such as Hayes, use a proprietary compression scheme during transmission that depends on hardware built into their modem. Because both sender and receiver must have a suitable modem, proprietary compression is unlikely to become popular.

- Multiplexers. A multiplexer (MUX) takes the incoming data streams from two or more computers and puts them on a single telephone line. Often an MUX includes a built-in data concentrator. Both sides of a communication link must have matching MUXs. Prices start at $1,200. A combination MUX and concentrator can run two channels at effectively 2400 bits per second with a single 1200-bps synchronous modem. Concentrators and MUXs work at the transport layer; many can handle only text.

COMMUNICATING WITH LARGE COMPUTERS

Nearly all computers consist of a processing unit, a keyboard, and a display. In a microcomputer, these elements all fit in the same package. A mainframe computer—with its very fast processing, high-capacity disk drives, and many megabytes of memory—is too big to fit in a single package. Its processing unit fills big boxes installed inside air-conditioned rooms; its keyboard-display units, or terminals, may be scattered all over. The traditional, or dumb, terminal does not contain any processing power and can operate only when connected to its large computer host.

To a large computer, the Mac is just another dumb terminal. But since the Mac has processing and storage capabilities of its own, it can act as an intelligent terminal, manipulating information before sending it on to the mainframe and storing the mainframe's responses for later use. For example, you can use most Macintosh word processors to prepare text before sending it to a mainframe. You can also feed information from a mainframe into a spreadsheet for further analysis.

Terminals come in many types, and suitable software will turn the Macintosh into most of the popular ones. The differences among terminals depend on many details, such as the code to erase a line on the screen. Although some terminal configurations truly offer special advantages, most came about because of arbitrary design decisions. Many communication programs emulate Digital Equipment Corporation's VT52 and VT100 terminals. Other companies have programs to emulate the DEC VT240, Tektronix graphics terminals, and many others.

With so many terminal types, some systems fall back on the simplest one: an electronic version of the old mechanical teletype. A teletype has minimal features—only the ability to print along a line and advance the paper; it won't even erase. In a rare display of nostalgia, many flashy terminals, including the Mac, can emulate a teletype.

When you buy software that makes the Macintosh emulate a terminal, be sure it emulates an intelligent terminal rather than a dumb one. Look for several key features:

- Disk storage and recording. The program should let you create information with Mac software, store it on disk, and then transfer the disk file to the mainframe; it should also let you store returning information on disk.
- Ability to store communication parameters on disk so that you can set them once and recall them as needed.
- Ability to suspend communications without disconnecting from the remote computer. This way, you can use a Macintosh program for a task and return to the other computer without having to reestablish a connection.

Synchronous Communications

Earlier in this chapter, we noted that synchronous communication works by sending information a block at a time, along with timing (synchronizing) information so that the receiving computer can extract characters from the block. A block of data is also called a frame or a packet. Telecommunication processes that deal with blocks of data, rather than voice or asynchronous data, are called packet-switching systems. (Voice can be managed if the packets are handled quickly enough.)

Because of substantial gains in transmission efficiency over asynchronous communication, mainframe computers usually communicate synchronously. For the best results with a microcomputer connected to a mainframe, the link should also be synchronous (even though some mainframes may have ports adapted to asynchronous operation). The synchronous link used by IBM and many other mainframes to talk with one another and to their terminals is known as Synchronous Data Link Control (SDLC), a version of the international standard High-Level Data Link Control (HDLC).

The physical link between the mainframe and the micro may be a wire or a modem. Unfortunately, many modems made for microcomputers cannot handle synchronous links. With suitable software and hardware, however, a Macintosh can be made to emulate synchronous terminals.

Several companies sell combination hardware-and-software protocol converters for the Macintosh, including Apple (AppleLine), Digital Communication Associates (Alpharetta, GA), and Avatar (Hopkinton, MA). Most of these emulate the IBM 3278 terminal. If several different Macs need to communicate only occasionally with a mainframe, a single converter can be switched among them. Some converters are cluster controllers, servicing several microcomputers or terminals simultaneously under software control.

Synchronous and asynchronous communications are usually coded differently. Virtually all microcomputers, including the IBM PCs, understand ASCII, but many large computers, particularly IBM and IBM copies, instead use EBCDIC (Expanded Binary

Coded Decimal Interchange Code). Conversion software simply looks up each incoming character in a table and sends out the translated code. The supporting software for all asynchronous-to-synchronous protocol converters will convert ASCII to EBCDIC.

FACSIMILE

In the corporate world, facsimile machines are common today for sending text and graphics over telephone lines. The widespread international Group 3 standard established by CCITT has two resolutions. Both have 203 dpi horizontally; standard mode has 98 dpi vertically and fine mode, 198 dpi. The Group 3 protocol includes a polling feature whereby the receiving fax machine can call up the sending fax to ask for transmission. In this way, a fax machine in a central office can check machines at all the branch offices. A password feature restricts document transmission to previously identified machines only.

Group 3 uses a compressed transmission format based on a modified version of V.29 half-duplex 9600-bps modems with automatic fall-back to 7200, 4800, and 2400 bps for noisy phone lines. To save on modem costs, companies such as Asher Technologies (Roswell, GA) have brought out fax modems that run no faster than 4800 bps. When greeted with one of these at the other end of the line, a standard 9600-bps fax modem will fall back to the slower speed. With suitable software, a fax modem can send not only pictures but also data to another fax modem. Fax modems are not, however, compatible with standard V.29 modems.

The Group 3 fax protocol was designed for traditional fax machines with built-in printers. Images are printed as they are received. While receiving, a fax machine cannot ask the sending machine to pause or stop; it can only hang up the phone line. Thus, there is no provision for a microcomputer-based fax system to pause and store information on a disk; fax software and hardware must receive and store the transmission simultaneously. Sending an image is easier than receiving one; a sending fax machine can pause briefly during transmission to retrieve information from a disk.

A new fax standard, Group 4, supports several resolutions, including 400 dpi— a major improvement over Group 3's 200 dpi. Group 4 is not yet widespread because its current standards require an ISDN phone line; CCITT has not yet specified a standard way for Group 4 to be sent over dial-up lines. The fax manufacturers that claim Group 4 operation over dial-up phone lines in early 1988 have adopted proprietary and mostly incompatible techniques.

20: Local Area Networks and AppleTalk

Personal computers were developed for individuals; they were bought individually and used individually. Communication with other individuals had to be done by modem or by swapping disks. Yet as micros gained power and became more common in companies, users began wanting to share information more readily.

The first practical way to share information among micros was to connect them as terminals to a larger minicomputer or mainframe, which could store files that the micros could access. But not everyone had a large computer; besides, all that expensive computing power could be put to much better use than shuttling files among micros.

Enter the local area network (LAN). Networks consist of nodes connected together. A node can be a computer, a printer, a disk drive, or any other device able to send or receive information or both. In most networks, any node can send a message to any other node; nodes not involved with a particular message ignore it. The computers in a network generally have equal access to network devices, which means they can share disk storage, printers, other resources, and electronic mail.

Early developers envisioned networks of diskless microcomputers sharing a central hard disk drive. Because no computer would need its own drive, costs were supposed to be lower. Things have turned out differently, however; putting a hard disk drive on every micro can now be cheaper than the combined costs of installing network wiring and a shared hard disk.

In some ways, networks are a step backward, for they represent a return to central control. With a freestanding microcomputer, you determine your own destiny. A network, on the other hand, must have a system administrator who sets up and manages it—and decides when to shut down a shared device. A breakdown of a critical shared component, such as a disk drive, can paralyze many users, a key problem with central computers. A network needs careful system management, including frequent file backups, to minimize such threats.

NETWORKS IN GENERAL

Networks come in a tremendous range of capabilities and prices. Before discussing AppleTalk, let's look at some principles of networks in general.

Network Topology

Bus networks have a single cable connecting each node. Each node has a unique address and monitors the network for information addressed to it. In the most common designs, each connection to the network is a passive circuit to ensure that a node failure does not disrupt the rest of the network. Adding nodes is simply a matter of tapping into the connecting cable. Most commercial networks, including AppleTalk, EtherNet, and WangNet, use this layout.

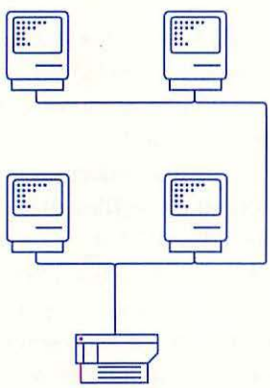

Bus network.

Ring networks are arranged around a closed loop cable. An information packet—analogous to a mail pouch—travels continuously around the loop. A repeater for each node relays the packet around the ring. When a node wants to send a message, it examines the packet. If the packet is carrying information, the node waits until it is empty. If the packet is empty, the node fills it with data, adds the address of the destination node, and sends the packet on its way. Each node looks at the packet for data sent to its address. When a node finds something addressed to it, it removes the information and sends on the empty packet. Ring networks can run over great distances because the signal is regenerated by each repeater, but if one repeater fails, the network fails as well.

Adding a node requires breaking the connection between two nodes and inserting a repeater. Ring networks are used in some universities and for IBM's token-ring network.

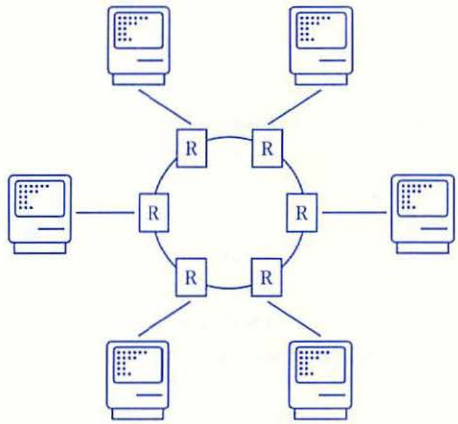

Ring network with repeaters (R).

Star networks have a central hub that handles all traffic. A mainframe computer serving multiple terminals acts like the hub of a star network. The hub is critical; if it fails, the entire network stops. Wiring costs are usually high because every node must be connected directly to the hub. A network can also use the private branch exchange (PBX) telephone system within an organization. PBX networks follow a star layout; the central switch becomes the hub. Wiring costs for PBX networks can be considered relatively low, since every office is wired for a telephone anyway. StarLAN from AT&T is a PBX-based star network.

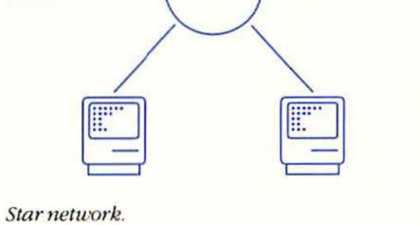

Star network.

Information Flow

As it travels, network information can be separated by time, by frequency, or by cable. (The speeds quoted here are raw data transfer rates, which are often misleading; see the discussion on page 313.)

A baseband network carries a single channel of information at a time. If the network is busy, each node waits for the channel to clear before sending information; messages are separated by time. Peak speeds range from slow to medium—about 50 kilobits to 20 megabits per second. AppleTalk and EtherNet are baseband networks.

A broadband network can carry multiple independent channels, much like cable television; messages are separated by frequency. A channel can carry data, voice, or even video. Broadband can be high speed, well over 20 megabits per second, and is generally expensive; both the cabling and the network interfaces cost more than baseband designs. The most elaborate versions use optical fibers as the communication pathway. WangNet is a broadband network.

PBX-based networks can be either analog or digital. Those based on older PBXs must use analog signaling and are slow, typically limited to 9600 bps. The new digital PBXs usually carry information at 64 Kbps, the data rate of an ISDN voice channel. Because a PBX network handles many messages simultaneously, the aggregate communication rate can be high. Within a particular cable, the signal is usually baseband, but separating messages is unnecessary because the messages are carried on separate cables.

Compatibility within a Network

Two nodes on a network can talk to each other only if their data can be mutually understood. Basic network communication protocol usually goes only through the network layer, but some networks have failed to achieve even this degree of compatibility because the various network hardware and software vendors do not agree. For useful information exchange, software must provide compatibility through the presentation and application layers. (See Chapter 19 for a discussion of protocol layers and Chapter 21 for some specific cases.)

Bridges between Networks

Most networks can handle only a limited number of nodes and a limited length of connecting cable. For efficiency, networks should usually have far fewer nodes than their maximum capacity; performance can suffer greatly when traffic approaches saturation. The maximum length of a single network cable is often set by the cable's electrical properties. Networks that use repeaters may not have a distance limit because the signal is regenerated at each node.

To circumvent the limits, a network bridge can connect two or more networks, each referred to as a zone. A bridge is a node on both networks that accepts messages from one zone and relays them to the second zone. Bridges must usually provide addressing services for the nodes in one zone to address those in the other. A suitably designed bridge can connect two local area networks that are thousands of miles apart.

A bridge that connects two dissimilar networks, converting protocols when required, is called a gateway.

A backbone network is generally a high-performance network used for high-speed communications and gateway services between two or more lower-speed networks.

Network Speed and Efficiency

The most often quoted speed figure is the raw data transfer rate, which is the fastest speed that bits can move along network cabling. But information actually travels at this rate only in brief bursts. In real situations, the overhead associated with a message—the address, conversion of files into frames, message confirmation, and so on—takes up considerable time. On bus networks there is always dead time, even when several nodes want to send messages.

True throughput speed—the rate at which useful information passes from node to node—is much less than the raw transfer rate, often as little as 20 or 30 percent; thus a "10-megabit-per-second" network may functionally move only 2 or 3 megabits per second. AppleTalk has a nominal speed of 230 Kbps, but its throughput speed is only about 65 Kbps with two active nodes.

Some network designs become unstable with very heavy traffic; they can even come to a complete stop. Careful planning is essential for a network that must carry heavy traffic.

Messages and Speed

The common fixation on the raw data transfer rate obscures the fact that speed doesn't matter for many network applications. Fast networks are expensive, and many offices don't really do anything that would justify such an expense.

Activities where speed is usually unimportant include the following:

- Electronic mail. Network speed is almost irrelevant; in most situations, it makes little difference whether the message gets to its destination in 0.1 second, 10 seconds, or even 10 minutes. For truly urgent messages, you would probably use the phone.
- Sharing files. Shared data files are most commonly treated like electronic mail, except that complete data files, instead of text messages, are sent. Unless a file is unusually large, it will get to the recipient within a few seconds regardless of network speed.
- Printing. Throughput is more often limited by printer speed than by network speed, unless extremely complex graphics are involved.

Some network traffic is moderately speed sensitive:

- Fully downloaded programs and data files. If you work with software and data that reside entirely in RAM, network speed is not always critical. You will have to wait for the network only when opening and saving files; the application does not need to go back to the disk while it is running.

Some traffic is highly speed sensitive:

- Disk-intensive programs and data files. If software running on a networked machine must frequently go to the disk drive to read in more program code or data, network speed can be a factor, especially if heavy traffic forces additional delays.
- Multiuser interactive software. Some new software will let two or more users work on the same file simultaneously; several people might edit a memo or draw a diagram together. Everyone sees the same screen (although only one person at a time can make changes). Because all screen updates are sent via the network, rapid response helps greatly. Few programs like this have appeared so far.
- Multiuser databases. When many people must have quick access to a complex shared database, network speed can be crucial. (Airline reservation systems are an extreme example, although they do not use standard network architecture.)

When a Network Is Overloaded

There are three general solutions for an overburdened network:

- Break the network into two smaller ones and reassign the most intensive users so that traffic on each network is cut down; use a bridge to maintain connections among all nodes.
- Give intensive users their own hard disks if they are taking up too much network time working with unshared files on a shared disk; a SCSI hard disk works faster than most networks.
- Reschedule the workload so that predictable periods of heavy use do not overlap.

NETWORK SOFTWARE COMPATIBILITY

Software designed for single-user computers often requires modification to run satisfactorily on a network. Some programs cannot run at all on a network because they cannot be started from a file server. Single-launch programs are those that can run from a file server, but only one person can use them at a time; additional copies are required for additional users. Multilaunch programs are those that permit several people to use them at once. Programs that can be configured to set fonts or other user preferences must be able to manage such information for each person on the network.

Programs intended for networks often feature file locking. Ordinarily, only one person can use a file at any given time; everyone else is locked out. In a variation of file locking, the first person to open a file can make changes in it; subsequent users can read the file but cannot change it.

Database programs designed for networks may have record locking, where several people can open a file, but only one person at a time can change a specific record. Once the record has been changed, other people can see the changed record, but no one can change it again until the original user releases it.

APPLETALK

AppleTalk is a relatively simple, low-cost bus network protocol that can tie together as many as 32 nodes. The raw data transfer rate is 230.4 kilobits (29 kilobytes) per second; actual throughput is about 65 kilobits (8 kilobytes) per second.

The physical connection between nodes consists of a small box containing an isolation transformer (to protect each node from electrical interference) and plugs for the network cables. An AppleTalk network uses a shielded 78-ohm twisted-pair cable, available in long rolls as well as short lengths (Belden 9272 or equivalent). Apple used to call this cabling AppleTalk cable but now refers to it as LocalTalk cable, to distinguish it from the AppleTalk network and protocol. The maximum length of an AppleTalk network is 1000 feet (300 meters).

When installing cable, it is a good idea to install a LocalTalk connector at every likely location for a computer or other node, even if you will not use it immediately. Otherwise, you will have to cut the cable to add nodes later.

An AppleTalk network can operate with ordinary twisted-pair telephone cable with an adapter from Farallon Computing (Berkeley, CA). With phone cable, maximum network length can be 3000 feet; for longer networks, Farallon has repeater boxes to boost the signal. Farallon also makes a star-configuration controller that organizes an AppleTalk network into a star topology; it is designed to work alongside a PBX.

DuPont manufactures an optical fiber connector for AppleTalk that makes it possible to lengthen the cable between nodes to 6900 feet.

With hardware that provides timing signals to each node, an AppleTalk network can operate at peak speeds greater than 230.4 Kbps. The AppleTalk interface within the Macintosh, however, is designed for 230.4 Kbps, so further interface hardware is also necessary to run that fast.

Node Types

AppleTalk supports a number of different node types.

- Computers. AppleTalk can handle a mix of computers; Macintoshes, IBM PCs, and others can all communicate on the same network. Apple, TOPS (Berkeley, CA), and Tangent Technologies (Norcross, GA) all offer AppleTalk interface cards for the IBM PC. As long as they have suitable software added to their own operating systems, all the computers can share network resources, such as printers and disk drives. Of course, shared information must be understood by all the networked computers; simple text files will be readable by all, but an IBM PC will not be able to do anything with a MacPaint file unless it has a compatible graphics program. Eventually, sophisticated network software will convert files from dissimilar computers automatically. (See Chapter 21.)

In network jargon, a device available to everyone on the network has the word *server* added to its name. Thus, a *printer* is available only to an individual user, but a *printer server* is available to all. A disk drive on a network can be:

- A disk server. This appears as a disk drive to network users, who can examine directories and read and write information on the disk. Often, disk servers use software that divides a single physical disk into several logical volumes, which behave as if they were individual disks. In some designs, only one person at a time has access to a particular logical volume; in others, several people can use a volume, but only one person at a time has access to a particular file.
- A file server. This does everything a disk server does and more. Not only can several people use the same logical volume, but, with multiuser software, they can use the same file simultaneously. File servers have their own processor and controlling software. An ordinary Macintosh with a hard disk drive, for example, can be turned into a file server.

  File servers can be dedicated or distributed. A networked Macintosh that does nothing but hold files is a dedicated file server; a given network can have one or more of these dedicated Macs. While operating as a file server, a Mac cannot run ordinary application programs.

  A distributed file server can run normal application programs in the foreground as well as serve files to the network in the background; any disk drive—floppy or hard—on any computer connected to the network can supply files to anyone on the network. If an application program requires a lot of attention from the file server's processor and disk drive, file serving may not work very well; when this happens, you can simply quit the foreground program. Foreground programs can cause more serious problems if they crash; usually the file server crashes as well, affecting more than one user.

  A file server can be another type of computer altogether (such as an IBM PC), or it can be built into a shared printer or other peripheral. Some file servers are specially designed computers and disk drive(s) without keyboards or screens. A file server's operating system and other details may be hidden from users, who see only normal Macintosh operation. Some file servers may have an operating system (such as UNIX) different from that of the Macs connected to them.

Other peripherals can also work as network "servers":

- Printer server. A printer on a network is normally available to all users, but if the printer is busy, anyone wanting to print has to wait. To eliminate the waiting, you can install several printer servers. Apple's printing resources can handle several LaserWriters at a time; you select the one you want to use. Alternatively, network software can store or spool printer information in a temporary disk file and then automatically forward the data to the printer when it is free.
- Mainframe server. Microcomputers on the network can operate as intelligent terminals to a mainframe computer, which performs protocol conversions and can talk with several terminals simultaneously.
- Modem server. A networked modem can either be connected to a computer on the network or operate as a network node in its own right. If connected to a computer, the modem can be a standard kind that you can manage with suitable network software. If the modem is a network node, it must contain its own network interface hardware and will thus cost more. Either way, only one person at a time can use a modem server.

  Should you get a modem server or a separate modem for each user? Multiple modems are more convenient if many people on the network frequently need one. Nowadays, modems are fairly cheap, although each modem may need its own telephone line. If modems are not used very often, and traffic on the network is not too heavy, then a modem server, which needs only one telephone line, may suffice. To support a modem, the network communication protocol is modified to send short blocks of information because short blocks take less network time to transfer.
- Port servers and peripherals. A node can supply standard serial and parallel ports to connect peripherals that are normally connected to a single-user computer. This arrangement can be practical if the peripheral sends and receives information relatively slowly. In contrast, peripherals that generate information rapidly, such as scanners, work best when they have the full attention of a host computer. They should be connected to a single computer that can forward the results to the network.
- Bridges. A bridge hooks together two similar networks. A local bridge, such as Interbridge from Hayes (Norcross, GA), connects two or more nearby AppleTalk networks. Stringing several networks together enables you to connect more than the 32 users AppleTalk can support on one network. A half-bridge connects remote AppleTalk systems; the link between networks might be a telephone line with modems. A backbone bridge connects several AppleTalk networks to a high-performance network; it combines the features of a local bridge and a gateway.
- Gateways. A gateway connects dissimilar networks. For example, FastPath from Kinetics (Walnut Creek, CA) connects AppleTalk and EtherNet networks.

File-serving Software

The three major programs are AppleShare from Apple, MacServe from InfoSphere (Portland, OR), and TOPS from TOPS (Berkeley, CA). All work adequately for standard network operations, and all support IBM PCs on the network. AppleShare is the least flexible because it turns a Macintosh into a dedicated file server, whereas MacServe and TOPS set up distributed file servers. TOPS leads in flexibility, because it can also act as a terminal for UNIX computers that use the TCI/IP protocol and Ungermann-Bass Net/One networks. MacServe is by far the least expensive; it is a good choice for networks that have only Macs.

Network Addressing

Going from the highest level to the lowest, an AppleTalk address can have as many as three levels: the network, the node, and the socket.

A network address is needed only when several networks are connected, in which case each network has its own address; AppleTalk can address up to 65,000 networks.

Within a network, each node has its own unique address. Some nodes have an address stored in nonerasable memory or on disk. If the node does not have a permanent address, it gives itself a random address number when turned on. To ensure against duplicate numbers, the node then sends a message to its own address; if any other node accepts the message, the sending node changes its number.

Users do not normally need to know node numbers; they deal with names—*Jane Doe* or *LaserWriter C*. When you first connect to the network, you give the AppleTalk software your name to store on disk. Thereafter, whenever you start a program that uses AppleTalk, your name is retrieved from the disk. When you send electronic mail, the network software polls all attached nodes and shows you a list of active users. You send your message by name; the software automatically attaches the correct node number to the message.

A socket is a logical concept, not a physical entity. A node can contain multiple sockets. Sockets generally specify particular application programs within the receiving microcomputer, so a message can be sent to someone's electronic mail program and not the database, for example. With socket addresses, the node can distinguish an electronic mail message sent by another user from a data file sent by a file server.

Network Control and Protocols

Once connected to a network, how does a node know when it can talk? On AppleTalk, EtherNet, and many other bus networks, there is no master controller; instead, every node can initiate a message on its own. The control comes from Carrier Sense Multiple Access with Collision Detection, or CSMA/CD. This mouthful describes a system with two simple rules:

- Any node currently using the network has priority. When a node needs to send information, it first looks at the network. If it senses a carrier, indicating that the network is in use, it waits. If the network is not in use, the node sends its message.
- If two nodes start transmitting at the same time, their messages collide; both nodes then back off and wait a random interval before trying again. Because of the independent delays, a collision is unlikely on the second try.

If the network is free, the sending node first transmits a Request to Send addressed to the receiving node. The receiver must acknowledge with Clear to Send within 200 milliseconds. If there is no acknowledgment, the sender must try again later. Some network software may deal with any busy nodes as a printer spooler would, automatically storing the message in a temporary disk file and forwarding it later without further action on the user's part. A message can also be broadcast to all nodes on the network.

Frames

Information sent on an AppleTalk network travels in frames organized in the manner of mainframe-to-terminal SDLC/HDLC links; each frame consists of several components:

- A synchronizing pulse marks the beginning of a frame; it is followed by two or more flag bytes marking the beginning of data.
- The frame preamble contains the addresses of the recipient and source nodes and the type of information contained in the frame.
- The data follow, as many as 600 bytes. The length of the data stream is variable; the receiving node determines the length from the synchronizing pulses (alternatively, software can set a fixed length). Some frames, such as the Request to Send, do not contain data.
- The frame trailer contains the Frame Check Sequence; sender and receiver independently calculate an error-correction value from the data in the frame. If the two calculated values do not match, the frame is discarded and the receiver requests that the frame be sent again.

Many messages are longer than 600 bytes, of course. As many as 12 frames can be linked together by a special identification code. If there are errors, the receiver need ask only for the faulty frames. Messages longer than 7200 bytes (12 × 600) must be broken into groups of frames, thus preventing two nodes from hogging the network. In addition, the receiver needs only a modestly sized memory buffer for incoming information. AppleTalk can operate in the background so your work is not disturbed when someone sends you a short message; long ones may require some disruptive disk activity.

ELECTRONIC MAIL

Sending a message from one user to another on a network is easy and doesn't require any special software. The file server keeps a mail folder for each user; to send messages, you merely create files and put them in recipients' folders. In groups of up to about 20 people, this method usually works adequately for electronic mail.

Specific electronic mail software adds many convenience features—organizing messages, letting you know when or if a message has been read, maintaining mailing lists so that you can send messages to individuals or groups, and so on. A flashing icon or beep can alert you whenever you have mail. InBox (Think Technologies, Lexington, MA) and Microsoft Mail are competing products that effectively and unobtrusively manage mail over AppleTalk networks. Both programs can extend mail services to IBM PCs attached to the network.

ETHERNET

AppleTalk is a low-cost, modest-performance network. If you need one that operates faster, the principal alternative is EtherNet, a bus network that uses CSMA/CD in much the same fashion as AppleTalk. EtherNet costs considerably more than AppleTalk—more than $700 per connection plus wiring—but it runs 20 to 40 times faster. Its nominal speed is 10 Mbps; actual throughput is 2 to 3 Mbps. Apple sells an EtherNet board (called EtherTalk) for the Mac II; Kinetics makes two EtherNet interface boards, one for Mac SEs and a second that plugs into the SCSI port.

Cabling for EtherNet can be thin coaxial R/G-58A/U cable, which is 0.2 inches in diameter and good for runs up to 300 meters (984 feet); thick coaxial R/G-11 cable, which is 0.4 inches in diameter and good for runs up to 500 meters (1640 feet); or twisted-pair telephone wiring.

Compatibility of Mac EtherNet products with other EtherNet hardware and software depends on specific configurations; unfortunately, EtherNet products are often incompatible. 3Com (Santa Clara, CA), a major EtherNet vendor in the IBM PC world, also makes file servers compatible with AppleTalk that can function as gateways to EtherNet-equipped IBM PCs. The major competitor of 3Com for IBM PC networks is Novell (Orem, UT), maker of NetWare. NetWare runs on a variety of networks, but usually on EtherNet cable. In 1988, Novell plans to market Mac products compatible with NetWare.

21: Moving Information Between Computers

Moving information between dissimilar computers involves two separate steps. First, you must establish a physical communication link between the computers; second, you must convert the data stored in one computer into a form understood by software running on the other. In this chapter, we emphasize moving information from another computer, mostly an IBM PC or clone, to a Mac. Usually, however, information can also be moved from a Mac to another type of computer.

PHYSICAL CONNECTION

You can link the computers physically by:

- Reading and writing 3½-inch MS-DOS floppy disks with Apple's upgraded floppy disk drives. (See Chapter 6.)
- Reading and writing 5¼-inch MS-DOS floppy disks with an accessory floppy disk drive for the Mac.
- Reading and writing Macintosh floppy disks directly on an IBM PC or clone.
- Connecting the computers via their serial ports with a null-modem cable and communication software on both machines. (See Chapters 13 and 19.)
- Establishing telecommunications via phone lines and modems instead of a null-modem cable. (See Chapters 13 and 19.)
- Establishing telecommunications through a third computer, such as an electronic mail service. (See Chapters 13 and 19.)

- Using the Mac as a terminal to another computer (terminal emulation). (See Chapters 13 and 19.)
- Connecting the computers via local area networks, most likely either AppleTalk or EtherNet. (See Chapter 20.)
- Using an OCR (optical character reader) to transfer text. (See the discussion in this chapter.)
- Using a disk converter to read and write IBM and Mac disks. (See the discussion in this chapter.)
- Sending the disks to a service bureau for conversion. (See the discussion in this chapter.)

The easiest way to transfer information is to copy files from one disk to another, just as if you were copying between Mac disks. From the hardware standpoint, copying is simplest when the MS-DOS disks you are transferring from are 3½-inch microfloppies and you have high-density (1.6 MB) disk drives on your Mac; all you need is file-transfer software and no additional hardware. Transferring from 5¼-inch MS-DOS disks is more complicated because your Mac will need an accessory floppy disk drive to read and write these. The Mac does not need an MS-DOS processor board to read MS-DOS disks, but if you have one, some file transfers are easier because you can use an MS-DOS program to modify the data before making the transfer. Apple II 3½-inch floppy disks can be dealt with in the same way. (For more information, see Chapter 6.)

Instead of using a Mac to read an MS-DOS disk, you can also do the reverse, reading and writing Mac disks with an IBM PC or clone. To do this, you need a Macintosh external floppy disk drive and the MatchMaker disk adapter card and software from MicroSolution (DeKalb, IL).

To move a file via telecommunications, you need to link the two computers, either with a null-modem cable or a pair of modems plugged into the telephone line. You also need communication programs for both the Mac and the other computer, preferably ones supporting a compatible data transfer format that can send binary files—those with any combination of bits, not only text. With ordinary communication programs, including those able to send binary files, you will need to operate both computers while transferring files. More specialized programs, such as MacLink Plus (DataViz, Norwalk, CT), are easier to use for moving files from an IBM PC or clone. MacLink Plus comes in two parts: one for the Mac and the other for the PC. Once the program is started on the PC, you can control everything from the Mac.

For some computers, you might not be able to find a communication program that supports a binary file-transfer format compatible with any Mac communication program. In this case, you will only be able to transfer text files, or you can use an electronic mail service—which can also be restricted to handling text files—as the intermediary. A binary file can, however, be converted to what looks like a text file and then converted back by software on the receiving computer. For example, Desktop Express (Dow Jones, Princeton, NJ) on the Mac can send binary files to Lotus Express (Lotus Development, Cambridge, MA) on an IBM PC via MCI Mail.

Terminal emulation programs set up a Mac to act as a terminal to a wide range of host computers. Most of these programs let you store any text sent to the screen by the host computer, so you can capture the text in a file. You might also be able to capture screen graphics as a bit-mapped image. Terminal emulation programs might not let you transfer a binary file from the host, however.

Local area networks almost always let you transfer files between any two computers on the network. (In some cases, you may need to move the file to a file server as an intermediate step.)

Optical character readers (OCRs), disk converters, and conversion services usually deal with text or word-processing files. (See page 327.)

To move files from Tandy model 100, 102, or 200 portable computers, you can use virtually any Macintosh communication program and a null-modem cable to the Tandy portable, all of which have built-in communication software. If you have Tandy's microfloppy disk drive for these portables, you can use MacDOS (Traveling Software, Bothell, WA) to plug the disk drive into a Macintosh and read and write disks in Tandy's format.

To exchange files with a DEC VAX, you can choose from several products. White Pine Software (Amherst, NH) does it simply and inexpensively by running DEC VT240 terminal emulator software on the Mac and file-management software on the VAX. Other links are much more expensive. Pacer Software (Westborough, MA) connects to VAXes through EtherNet and an AppleTalk gateway from Kinetics (Walnut Creek, CA). Alisa Systems (Pasadena, CA) sells a comprehensive family of Mac-to-VAX networking products called AlisaTalk. The Alisa product family includes Apple's AppleTalk software for the VAX VMS operating system.

LOGICAL CONVERSION

After you have physically moved a file, what can—or must—you do with it? If a Mac application can open and understand the foreign file format immediately, you won't have to do anything before using the information. More likely, you will have to perform a second step, logical conversion, to ensure that the transferred file is intelligible to a Mac application.

Logical conversion involves changing portions of the file, or "file massaging." Sometimes file massaging is built into the transfer hardware. What you need to do depends on what kind of file you have moved.

Word-processing Files

Most programs that create a text document store two kinds of information: the text itself and formatting—left and right margins, page breaks, tabs, headers, and so on. Formatting information is generally interleaved with the text; no standard formatting system exists, not even in the Macintosh world.

What you can transfer depends on the kind of text information you have. The best way to work with a foreign word-processing file is by reading the original, or native, format directly or using a conversion program. When this is impossible, the next best choice is to find a revisable document format.

Native document

This term refers to the file normally created by a word processor, complete with formatting information. It is easiest to convert native files if you have a Mac word processor that can read and write files created by an MS-DOS word processor. Microsoft Word and WordPerfect for the Mac, for example, can read and write files from their respective MS-DOS namesakes.

Utility programs that can convert one file format to another are another easy solution, although the quality of format conversion varies widely. Most of the conversion programs available so far have been for the MS-DOS market; they convert the files of one MS-DOS word processor to the format of another. Using such a program (for example, Word for Word from Mastersoft, Phoenix, AZ), you can convert files from a variety of MS-DOS word processors to the MS-DOS version of Microsoft Word or WordPerfect and then read these files with the corresponding Mac word processor.

Revisable document

In a revisable document, the formatting in the original file has been converted into a form that you can edit with the receiving computer and software. Although the formatting information might be incomplete, essential parameters such as left and right margins, tabs, columns, and line spacing have been preserved. The format coding on the converted document should be in the form of the receiving software. Software that converts to MacWrite's format should, for example, set the left margin by the ruler in MacWrite, not by putting spaces at the beginning of each line.

IBM's *DISOSS (Distributed Office Support System)* is a standard interchange format specified by IBM for its office computer products as a part of SNA (System Network Architecture). The overall strategy consists of two key format designs, or architectures:

- DCA (Document Content Architecture) defines such attributes as page width, tabs, headings, and so on. DCA itself comes in two forms: FFTDCA, or Final-Form-Text DCA, and RFTDCA, or Revisable-Form-Text DCA.

Final-Form-Text DCA:	*Revisable-Form-Text DCA:*
Top margin location	Declare top and bottom margins
Left margin location	Number pages and lines
Line spacing	Specify space occupied by body text
Font definition	Specify page width and height
Justify text	Insert fields from external data records
Begin and end underscore	Include text from other documents
Begin and end overstrike	Keep specified text together on the same page
	Spelling verification control

- DIA (Document Interchange Architecture) specifies the communication protocols for sending DCA files. DIA covers not only specific transmission structure but also how requests are made for a file. It can create a document library for multiple users. Although DIA specifications describe only text files, other file types can be included by extending the specification.

DISOSS, used as an intermediary between two otherwise incompatible systems, is becoming the *lingua franca* of office word processing. Many companies that produce word-processing software and hardware have announced support for at least part of DISOSS, usually DCA. Some companies support only FFTDCA because they fear that using RFTDCA may make it easy for their customers to change over to IBM or other RFTDCA-compatible products. Some independent software developers offer more complete DCA support than original word-processing hardware vendors do.

On the Macintosh, Microsoft Word, version 3.0, and PageMaker, version 2.0, can accept DCA files; more products will do so in the future. DCA does not support some standard Mac features, including different fonts, type sizes, and graphics. Most other office computers either do not understand such information or use it in incompatible ways. You need only be concerned about DCA support if you must communicate with a DCA-compatible computer system. The Macintosh with almost any word processor is actually much better for editing words than any traditional office word processor.

RTF (rich text format) was developed by Microsoft as a word-processing exchange format that improves on DCA by managing more text attributes and features. RTF is stored in pure ASCII and can be sent via any electronic mail service. The value of RTF will depend on how widely the computer industry supports it.

If your word processor cannot read a foreign file you need to transfer, and neither a conversion program nor a revisable format is available, you will have to make do without the original formatting. You can usually still move the file, but it may need considerable reworking with the receiving word processor before it is legible. For short files, this won't be a big problem, but reworking long files can be tedious.

Text-only document

Stripping the formatting out of a raw document leaves a text-only document, which is generally easier to work with. Most electronic mail systems accept text-only documents, although they may limit line length. Many microcomputer word-processing programs can create text-only files. Such files created on a Mac can contain non-ASCII characters (those created when you press the Option key, such as " " ' ') that are not understood by other computers; these may appear as spurious characters on the other machine.

Text-only documents created on MS-DOS computers usually have both an LF (linefeed) and a CR (carriage return) character at the end of each line. The Macintosh uses only CR characters at the end of each line, so if you read an MS-DOS text file, you will need to get rid of the LF characters, which will show up as open rectangles in your document. You can do this with utility programs available from user groups; some Mac word processors let you search for LF (its code in ASCII is Control-j; in decimal, 010; in hexadecimal, 0AH) and replace it with nothing.

Final-form or page-image document

You cannot always work with a text-only file. Some word processors, for instance, will not create them; others clutter the raw document with embedded formatting information that will take some effort to strip out. In such cases, you should try to transfer a page-image file. A page image is the disk equivalent of the printed page, complete with headers, footers, and page numbering. Many word processors on other computers let you make a page image with a command such as "print to disk." These word processors are designed for printing documents with a daisy-wheel printer and fixed character size and spacing, so the page image consists of only ASCII text. Macintosh word processors usually don't have a page-image function because they routinely deal with variable fonts and graphics.

Page images are not particularly easy to work with, but at least they are free of embedded formatting commands. They do, however, have a carriage return, and often linefeeds as well, at the end of every line. If you use a Mac to work on page-image files, you must remove these characters to restore wordwrap.

If your word-processing program permits, the method that follows is the quickest way to get rid of unwanted carriage returns. You want to preserve the carriage return at the end of each paragraph, so search for all returns followed by a tab (or five spaces, or whatever indicates paragraph indents) and replace them with a # or some other rarely used character or sequence of characters. Then replace every remaining carriage return with a space or with nothing, depending on whether the last word on each line is followed by a space before the word that begins the next line. Finally, replace the # symbol with a carriage return and your paragraph indent character(s). Although this procedure may mangle a few lines, it's faster than deleting carriage returns one by one.

A few word-processing programs—MicroPro's WordStar, for example—store text using eight bits for each character instead of the more common seven bits; the extra bit must be stripped off before moving files created by these programs. Mac user groups can supply public-domain software for doing this, such as UNWS or WS to MW (for converting to MacWrite document format).

Other Document-Conversion Options

It is relatively easy to shuttle information in and out of microcomputers; getting information out of other kinds of computers is another matter.

Dedicated word processors

Before the rise of microcomputers, many companies produced dedicated word processors, which are specialized microcomputers designed only for word processing. Wang word processors remain the most widespread of these machines. To move text files between a Mac and a Wang system, you can use the Wang's telecommunications option. On an asynchronous line, the Wang will send out text-only files. For a fully formatted file, you will need a bisynchronous protocol converter or a Wang-to-DCA conversion product. Omnigate (San Francisco, CA) sells Wang VS-to-Mac conversion items.

Disk converters

A dozen companies make disk converters: freestanding computers, each with a highly flexible disk controller and multiple disk drives in several sizes. The converter can read files in many different physical and logical formats and write them as well. Software within the converter converts from one format to another. A few companies have announced support for 3½-inch disks, but you must check specifically for Macintosh format. At the high end, an Altertext converter (Boston, MA) with a 3½-inch floppy option costs more than $20,000; low-end units such as one from Flagstaff Engineering (Flagstaff, AZ) can be as low as $3,000 plus a host IBM PC. The Keyword 7000 system (about $11,000; Calgary, Alberta) offers particularly complete text-file format conversion.

Optical Character Readers

If transferring a text file from another computer proves impossible, you can use an optical character reader (OCR) to scan a printed page; it "reads" the text and creates a text file for a word processor. Optical character recognition (also tagged OCR) is the only way to automate input from electric and most electronic typewriters.

OCRs have been finicky, unreliable machines; they can usually read only originals, although the best can read clean photocopies, and they can rarely read dot-matrix printer output or proportionally spaced text from a typesetter or laser printer. Most OCRs are limited to specific fonts, usually the popular monospaced typewriter fonts such as Courier. Early OCRs could read only special fonts such as OCR A and OCR B, but few are so limited now.

All OCRs make errors; how many errors you can tolerate depends on the material you are reading and your ability to type in corrections. Even a one-percent error rate is usually too high because a competent typist can type a page faster than it takes to search for errors. A spelling program can help out for some documents. To reduce errors, some OCR software can be set to recognize only numbers in a particular region of a page—

such as a table of numbers for input into a spreadsheet program. By restricting the recognition to numbers, the range of possible errors is reduced. Errors in numbers, however, are harder to spot than errors in text.

Low-cost OCR systems consist of a scanner with OCR software that runs as an application program on a Mac. Most companies that sell scanners are offering such software. Some service bureaus offer a document-conversion service with OCRs.

The best OCRs to date have used specialized computing hardware designed specifically for the purpose. Two expensive (more than $30,000) OCRs work fairly well: the Kurzweil 4000 (Kurzweil Computer Products, Cambridge, MA) and the Palantir Compound Document Processor (Palantir, Santa Clara, CA).

Service bureaus

Many service bureaus have OCR equipment and disk converters; they offer both media conversion (transferring files between physically incompatible formats) and text-file format conversion. But costs are fairly high, and turnaround is slow. Some companies offer format conversions over telephone lines, using remote computers to do the work.

Other Non-Graphics Files

Generally, each spreadsheet, database, or other program stores files in a unique way. If a program is very popular, however, its file format can become a de facto standard that other programs adopt or can at least read and write. For MS-DOS computers, Lotus 1-2-3 is such a program. Microsoft Excel, Modern Jazz, and MacCalc (Bravo Technologies, Gilroy, CA) can all read and write 1-2-3 files on the Macintosh. Foxbase/Mac (Fox Software, Perrysburg, OH) and McMax (Nantucket Software, Los Angeles, CA) are both clones of the MS-DOS version of dBASE developed by Ashton-Tate (Torrance, CA). Both programs can not only read and write dBASE files, but work the same way dBASE does. Ashton-Tate's dBASE Mac can also read MS-DOS dBASE files, although its operational style and file format differ from those of the MS-DOS version.

File Interchange Formats

If you cannot get one program to read another program's files directly, then you should see if both programs support a common file interchange format. Although several such formats exist, support for them has not been widespread.

The DIF format was created by Software Arts, the company that wrote VisiCalc (and is now absorbed by Lotus). DIF files work at the presentation protocol layer and contain only printable data; they do not store the formulas and relationships of VisiCalc or other programs. The DIF format has found its widest use in moving spreadsheet data to graphing programs.

The SYmbolic LinK (SYLK) format, devised by Microsoft, can store formulas and relationships; several companies including Lotus (in Modern Jazz) support this format.

SYLK is the only common microcomputer data exchange format that functions at the application protocol layer. Because SYLK was designed for universal application, it is somewhat large and inefficient and used only for transferring information.

Few programs use SYLK as their native file format; you must explicitly tell your program to generate a SYLK file. Once generated, data files in SYLK can be moved to any other computer that has a program capable of reading SYLK files; all original functions remain intact—to the extent that the second program can understand them.

Both DIF and SYLK files can be transferred as ordinary ASCII; they do not require special handling, although transport and session layer protocols will make transfers much easier.

ASCII files

Many programs can produce ASCII files on disk that are equivalent to printed output, enabling the information the programs contain to be added to a word-processing document. These ASCII files usually lack formulas and relationships.

For programs that contain information in ordered, discrete blocks, such as cells in spreadsheets or records and files in databases, the ASCII output can often be delimited with a comma or tab placed between each block. Without delimiting, the information is useful only for transferring to a word processor. With delimiting, the information can be placed, block by block, into another spreadsheet or database.

Tab-delimited files

One way to create delimiters in a spreadsheet program is to add a fixed number of spaces to each column—many more than needed for any cell. Save an ASCII representation of your file that includes these spaces; then search for and replace them with a tab. This may be a multistep process, for you must then remove excess spaces and take care that you have not inadvertently deleted empty cells.

You can write a program in BASIC or another programming language to read the original file and convert it to another form or to create a tab-delimited file.

Several spreadsheets and database programs for the Mac can read tab-delimited data. For spreadsheets, each block is placed in a cell, and each tab moves the following block of data into the following cell; a carriage return moves down to the first column of the next line. For databases, tabs mark the ends of fields; a carriage return marks the end of a record.

Comma-delimited files

Comma-delimited data pose several problems. Commas are often part of large figures (as in 1,000,000), and Europeans use them instead of decimal points in numbers. Because of this confusion, many programs do not read comma-delimited data. If you need to move such a data file, you must first convert the commas to tabs.

If there are no commas or decimal commas in numbers, this conversion is easy with a simple BASIC program, or it can be done on many word processors. Use the word processor's global search-and-replace function, but first be sure the program can replace with a tab (some programs can search for and replace only ordinary characters). If your program cannot search for and replace with tabs, see if it will recognize Control-i, the ASCII code for tab. You may need to change valid commas within fields to another character temporarily and restore them afterwards.

Unless you are moving a SYLK-format file, all this effort still leaves you with incomplete information because the formulas and relationships are lost in the transfer. Nevertheless, for a spreadsheet model, the headings and number entries will be correct, and reconstructing the formulas is simpler than starting from scratch.

Moving Graphics Files

Bit-mapped graphics of the kind MacPaint makes follow no standard, but many bit-mapped files use a relatively simple format, and programs have been written to convert MacPaint files so that they can be read by IBM PC, Atari ST, and Amiga computers. These programs, mostly free or low cost, are available from user groups and bulletin boards.

If you want to move screen graphics, and the originating computer uses a standard RS-170 video-signal format, you can capture the screen image with a video digitizer (see Chapter 12) and get a MacPaint file. Or you can use a scanner on the printed output.

For object-oriented graphics, the most widely used formats were developed for computer-aided design (CAD). These include IGES (Initial Graphics Exchange Specification) and PHIGS (Programmer's Hierarchical Interactive Graphics Standard), both developed for elaborate CAD systems. On the IBM PC, the leading CAD program is AutoCAD; several companies support its file format. Some Mac CAD software, such as Pegasys, can read and write these formats. (See Chapter 12 for information about other graphics formats common to the Macintosh and IBM PCs.)

Programs

The machine-language form of a program—the form actually processed by the computer—can be understood only by another computer of the same type, so there is no point in moving a Mac program except to another Macintosh. Moving programs from an MS-DOS machine to a Mac does make sense if you have an MS-DOS board for your Mac or a software MS-DOS emulator; you can try any of the techniques outlined above. If you use communication programs with modems or a null-modem cable, choose a transport layer protocol that handles binary files.

Source code for a program—the original instructions as written in BASIC, Pascal, and other languages—can be transferred as an ASCII file. (With Microsoft BASIC, you should save the program with SAVE "FILENAME",A.) Depending on the program's origin, some adjustments may be necessary to run it on a Mac. You will need an interpreter or compiler in the appropriate language to turn the source code into instructions for the Macintosh.

GETTING INFORMATION FROM A LARGE COMPUTER

All the problems of moving information from one microcomputer to another also apply to large computers. Mainframes are at their best handling large databases and crunching huge arrays of numbers. Because large computers are always managed by data-processing specialists, the software they use appeals to the specialist rather than to the ordinary user. Or, to put it another way, mainframe software is much harder to use than microcomputer software, especially Macintosh software.

Getting information out of a mainframe and into a micro isn't easy. In most cases, you must go through a normal session with the micro acting as a terminal to the mainframe database and store all the results on disk in the micro. Then you must spend considerable time manipulating the data into a form suitable for a microcomputer program. The whole process is so tedious that most people don't bother to do it electronically; instead, they print out the results and type them into the microcomputer. Then the results can be looked at with a spreadsheet or graphing program.

A few microcomputer programs can now perform this transfer with less pain. These programs act as buffers between you and the mainframe. You ask for information through a microcomputer interface; the program converts your requests into commands understood by the mainframe database manager. Information from the database then passes through the microcomputer and appears in a spreadsheet, ready for your next step. In time, all important mainframe programs will have such microcomputer support; until then, there will be much wasted and repetitive work.

A similar problem applies to many databases stored on mainframes but designed for access by microcomputers. All the popular dial-up information services currently use awkward, obsolescent interfaces that appeal only to determined computer hobbyists and specialists. Dow Jones, however, does have a special software product that hides some of the problems: Spreadsheet Link is specifically designed to extract financial information from the Dow Jones database and put it into the SYLK format. The program requires the Dow Jones Straight Talk communications program.

System Network Architecture and APPC

System Network Architecture (SNA) is the name IBM gives to a wide range of communication protocols and standards. Many early SNA standards, including some that define mainframe-to-terminal connections, are already widely used; others have yet to be specified. System Application Architecture (SAA) is one of these; it is supposed to be a unified software interface that makes an application look the same whether it is running on a mainframe or a small computer.

One key component of SNA is Advanced Program-to-Program Communications (APPC), a protocol for exchanging any kind of information between two or more computers. APPC protocol consists of a program-to-program portion called Logical Unit 6.2 (LU 6.2) and a network management portion called Physical Unit 2.1 (PU 2.1). It is not limited to simple file transfers; information can be exchanged between disparate programs and disparate machines. Each program on each machine must have its own version of APPC built in, however; there is no universal interface. Several APPC exchanges among different programs and computers can take place at the same time. Under APPC, a program on one computer can, for example, get a second computer to start a spreadsheet, load some financial data, perform a computation, and then forward the results to the first machine.

Conventional mainframe-to-terminal (or mainframe-to-microcomputer) connections with earlier SNA protocols are based on a master-slave relationship. Someone at the terminal or micro must log on, request attention from the host mainframe, and then sit there; the terminal or micro cannot operate on its own. In contrast, APPC is a peer-to-peer protocol; no one computer in the transaction has an inherently more important status than any other. A mainframe program can ask for information from a micro just as readily as the reverse. Furthermore, APPC can operate without human intervention, taking care of such problems as a breakdown in communication by reestablishing a link on its own.

Apple has announced a combination hardware and software package, MacAPPC, for the Mac. The LU 6.2 portion of MacAPPC is licensed from Orion Network Systems (Berkeley, CA). The hardware is an accessory circuit board, the Macintosh Co-processor Platform, which plugs into a Mac II. It has a 68000 CPU running at 10 MHz and 512 KB of RAM; PU 2.1 is built in. On a given AppleTalk network, only one Mac II must have the Co-processor Platform installed; other Macs on the network can communicate with APPC through that computer.

Apple's MacWorkStation software is a related product; it is a developer's toolkit for creating programs on a Mac that will interact with another computer using APPC.

22: Reproducing the Macintosh Screen

This chapter deals with reproducing Macintosh screens, primarily with photographic and video equipment—in case you wanted to distribute a video image for training purposes, for example, or if you were using the Mac as a prop in a photograph. Much of the chapter is highly technical; for each section, we assume that the reader has the background appropriate to the topic at hand. Essential terms are defined in the Glossary, but this chapter is no substitute for a technical handbook on photography or video. For the most part, you will get satisfactory results if you follow only as much of the advice as your equipment can handle.

AN OVERVIEW

Printing the screen image on paper as a screen dump is the simplest way to reproduce the Macintosh screen. The quality is low to medium, and the result is normally limited to black and white.

Typesetting the screen image gives a high-quality reproduction; this is also normally limited to black and white, although color is possible through color separations.

Photographing the screen directly can be done as:

- Still photography. The image will be low to medium quality; the process is inexpensive ($250) and quick.
- Motion picture photography. The images will be very low quality without special equipment and low quality with expensive equipment (more than $10,000).

Uneven refreshing will be less of a problem if you overexpose a little, but the key to even exposure is slow shutter speeds. You should get passable results at 1/8 second and better results with slower speeds. At a shutter speed of one second, the density variation will be at most 1/60 of the overall exposure, and invisible. Don't forget to use a tripod at such slow speeds.

Note that different CRTs vary in their light-decay characteristics, so experience with other computers or television sets may not apply to a particular Macintosh screen.

To set the aperture, take a through-the-lens exposure reading on any mostly white image, and then open the aperture two f/stops wider for your first trial; this will make the whites white, instead of medium gray. Standard light meters may not be accurate if their spectral sensitivity does not match the screen's spectral balance; readings may be off by as much as two f/stops. Bracket exposures by half stops the first time, three stops in each direction, and keep a record.

For most macro lenses, the best f/stops will be f/8 or f/11. Smaller apertures give less sharp results, and larger apertures make focusing and depth of field more critical.

Film

You can photograph a screen image with black-and-white or color film as a print or transparency. If you plan to project the photograph, you will most likely want a positive image (black letters on a white background, as on the Mac screen), especially because some graphics and icons are hard to understand when shown as negatives. The exception may be if you must project in a room that is not entirely dark, where negative images are easier to read.

Black-and-white positive films are rare; the only common high-contrast positive film, which is the best choice and also gives quick results, is Polagraph HC instant slide film for Polaroid's instant 35-mm Autoprocess system. The other two Polaroid instant slide films (continuous-tone black-and-white Polapan CT and color Polachrome CS) don't work as well.

If these films are unavailable or not suitable for some reason, you can check with graphic arts facilities to see which black-and-white processes they can handle, or you can process the film yourself. Most graphic arts labs use a two-step process, printing negative-to-negative to get a positive. Kodak High-Contrast Copy film works well for original black-and-white negatives; as with all high-contrast materials, the exposure latitude is small. With any black-and-white film, no filters are needed.

If you don't have a convenient graphic arts facility nearby, you can shoot color film and process it commercially. If you do, however, the Mac screen will appear bluish in your photos; the exact color will depend on your exposure and emulsion. This happens because the color balance of CRTs depends on the phosphor coating that creates the light. Small Mac screens have a P-4 two-component phosphor with two emission peaks, at 460 nanometers (nm) and 560 nm, in the blue and green portions of the spectrum.

You can see the two components with a magnifying glass; some regions are a little more bluish than others.

Although the concept of color temperature applies only to smooth emission curves, not the P-4 spectrum profile, the rough equivalent color temperature is about 11,000 degrees Kelvin.

Color Correction

Optimum correction of the bluish cast of small Mac screens requires some experimenting with your exposure and film emulsion preferences. To start, use daylight color film and a color-balancing filter. An 85C filter should yield neutral whites; the more common 85 filter also works but produces a warmer tone. You may need a CC (Color Correction) filter as well. Kenko's TV-CC filter may work effectively, depending on the film emulsion. If you must use tungsten-light color film, try to find the rare 86 filter, or combine an 85B with an 81EF. Kodachrome emulsions have better contrast and sharpness than other color films, but fewer laboratories can process the film.

The color balance of a Mac II depends on the video screen that is used with it; most—including color screens—are bluish, similar to the small Mac screens. To determine the balance, use a color temperature meter if one is available. If not, you can find the approximate temperature by placing a white card next to the Mac screen. Illuminate the card with a lamp or slide projector whose color temperature you know, and place neutral-density filters over the light source until the brightness of the card matches the brightness of the screen. Then insert color-balancing filters until the colors match. The color correction that you need is the opposite of the color-balancing filters.

Some photographers and art directors may question the need for color correction, arguing that if the screen is bluish, then it should look that way in print. Perhaps they think the blue feels cool and high-tech. When you look directly at the Macintosh screen, however, it appears white, not blue. The human visual system adjusts the predominant illumination to white in the same way that it adjusts to daylight or tungsten light.

False Color

With the appropriate filters over the camera lens, you can create color images from a monochrome Mac screen. You should prepare separate screen displays for each color you want, and then photograph them one at a time, changing the filter for each step of a multiple-exposure image. Color images are easier to make if the image on the screen is reversed so that text appears as white letters on a black background. If the background were white, the first exposure would effectively block the addition of colored objects on subsequent exposures. If your application program cannot reverse the screen, you can make a screen dump and then open it with Ashton-Tate's FullPaint. Within FullPaint, you can reverse the image and create a screen display that covers the entire screen, eliminating even the menu bar at the top. FullPaint, version 1.0, however, can only create displays the size of a small Mac screen.

Photographing the Screen as Part of a Larger Picture

If the Macintosh screen is only part of a picture, the bluish hue can look out of place. You can correct the blue by retouching the photograph or by choosing one of these two exposure strategies:

- Single long exposure, if all elements are static and the overall light level can be balanced against the screen.
- Double exposure, once for the scene and once with a long exposure for the screen alone. For the first exposure, black out the screen with photographic velvet; for the second, extinguish all other light sources, and use a matte box if necessary to reduce spilled light from the screen.

To balance the color, you can:

- Filter the ambient light to match the color balance of the Mac screen. If you have people in the scene, try an 80C or 80B gel over electronic flash units and a warming filter over the camera lens. Run tests if the results are critical; small errors can be corrected in the lab.
- Filter a double-exposed image separately, using the filters appropriate to the light source.
- Place a color gel over the screen. An 85 gel, available from motion picture supply houses, may be roughly correct.
- Take the screen photograph at a different time and insert it as a photo composite. Achieving high-quality results with this method is difficult, except with digital color-image processors, which are relatively easy to use but extremely expensive. (These image processors, such as the one made by Scitex, are used for jobs such as assembling the Sears catalog.)

MOTION PICTURE PHOTOGRAPHY

With ordinary cameras, motion picture photography of CRTs always shows bars moving rapidly across the screen because the video scan rate is not synchronized with the motion picture framing rate. If you have the resources, here's how to get the best possible results with a small Mac.

Get a motion picture camera with mirror reflex viewing, a 180-degree shutter, and an input for external speed control. For best results, the shutter should move vertically from bottom to top (the direction in which the Macintosh screen is refreshed, inverted for the camera lens).

To control the camera's speed, bring the 60.15-Hz vertical-sync signal out of a small Macintosh (the signal is on a wire connecting the CPU board and the video/power-supply board) and divide it by 2 (to 30.08 Hz); use the signal to control the camera's motor. (The division won't be necessary with some units.) With a good phase-locked system, you can set the controller so that the camera shutter is open as the image is written on the screen.

Alternatively, you can use a separate, adjustable precision oscillator for the camera control signal. The oscillator must be accurate to $1/1000$ of a frame to maintain synchronization; the longer the take, the higher the precision necessary. Be sure the shutter is set to 180 degrees, and run the camera without film to fine-tune the speed. For each take, you must adjust the shutter phase on the speed control: Hold down the phase button until the thin sync bar rolls out of the CRT image, and release it before the line returns. If the bar is out of the eyepiece with a 180-degree shutter, it is out of the film; the film will record every other frame of the Macintosh screen.

For synchronized sound, simply operate normally; the camera control signal will drive the Pilottone track on a Nagra or other recorder at a frequency 25 percent higher than normal, which poses no problems for most resolving equipment. (A Nagra IV will accept a 5-volt square wave with a 50 percent duty cycle for the pilot.)

After you have the film and sound track, you can transfer them to videotape: one frame of film to one frame (two fields) of video instead of the ordinary 24-frames-per-second (fps) film transfer. The audio should be resolved to a magnetic film recorder running at 30 fps. The slight speed shift downward to the broadcast-standard 59.94-Hz NTSC field rate isn't important.

If the result will be used in a 24-fps motion picture, then you will have to convert speeds with a step printer, essentially dropping every fifth frame; transfer the sound at 24 fps. Step printing is expensive and may result in minor motion discontinuities.

The method described above generally works better than the classic technique for filming a television screen, which does have the virtue of producing a 24-fps film. This technique requires that you use a movie camera with a 144-degree shutter and external sync; slave the camera motor to the Mac's vertical-sync signal.

The common practice of filming European television systems (50-Hz field rate) at 25 fps with a 180-degree shutter will not work with a small Macintosh.

In some situations, you may get better results by replacing the Mac CRT with another design whose phosphor persists longer. If you have the technical expertise, you can also change the Mac's clock speed and thus the rate at which the screen is refreshed.

Some Macintosh II video cards can synchronize their screen display signal to an external pulse generator, a feature known as genlocking (generator locking). This ability means that the screen can be synchronized to professional motion picture equipment in the same way that standard television images can be.

Step-Frame Motion Picture Photography

For best results, script and budget permitting, shoot the motion picture with double exposure. On the first pass, black out the Macintosh screen and shoot the live action or other surrounding material. Then shoot the screen on the second pass with a step-frame motor control, exposing each frame as a still image for $1/8$ second or longer with filtration. Use computerized stepping motors for camera movements if necessary. Matching the action may require a traveling matte shot. You could create the matte with a program that generates a blank white screen. For a clean matte with adjustable brightness, take an empty Macintosh case and replace the screen with a translucent white screen. Cut a mask that matches the Mac's screen area and illuminate the screen with a floodlight; use a filter for a blue screen matte.

ELECTRONIC USE OF THE MACINTOSH VIDEO SIGNAL

The video signal created by a Macintosh to drive its screen is called a native video signal; a standard video signal is one that is compatible with conventional video equipment developed for broadcast television. Native and standard video signals differ in many ways; one key difference is horizontal sweep frequency—the rate at which lines are written on the screen. Standard video equipment is designed to handle sweep frequencies of 15.738 kHz (North America and Japan) or 15.625 kHz (Europe); circuitry in such equipment cannot manage higher sweep frequencies.

Mac II video boards generate native video signals with much higher sweep frequencies. Apple's Macintosh II Video Card sweeps at 32.4 kHz; video boards designed for large screens can run at 70 kHz or more. Small Macs generate a native video signal that has a horizontal sweep frequency of 22.255 kHz. (Unlike Mac IIs, small Macs lack a video output connector, but installing such a connector is relatively simple.)

Native video signals will not work with ordinary television equipment, so what can you do with them? Some video equipment can cope with a Mac's native signal:

- A high-quality monitor sold for computers can display Mac images directly in monochrome and color.
- A projection monitor can create a large image suitable for lecture halls. The most widely used projection monitors generate the image with a CRT. A new class of projection monitor consists of a liquid crystal display that is placed in an overhead projector. These LCD models are much smaller and lighter than a CRT model, but the image is usually much poorer.
- A slide maker consists of a photographic camera and a video monitor in a specially designed lighttight box. In simple slide makers, the video monitor displays the same image as the computer screen; the slide maker is connected to the video output. More elaborate models have many more pixels than a Mac screen; they connect to Macs as if they were printers.

To send Mac images into ordinary video equipment, you must convert the Mac's native video signal to a standard form with slower sweep frequencies and lower resolution.

The simplest and crudest way to get a standard signal out of a Mac is to point a television camera at it. The results are poor, with a major loss of sharpness. Nevertheless, with a good camera, the results can be good enough for casual training and other less critical uses. Aside from the loss of sharpness, the difference in framing rates between a Mac and a 60-Hz or 59.94-Hz television produces a line or sync bar that runs slowly down the screen every few seconds. The line flutters more disturbingly on a European 50-Hz television.

For a small Mac, the only way to get rid of the bar is to slave the television system to the computer's vertical sync. In a studio, slave the studio sync line to the Macintosh and operate all equipment, including the videotape recorder, at 60.15 Hz.

As with other methods that combine two scanning processes, moiré problems may emerge at some magnifications. (If the Mac image is black and white, you can kill the color-burst signal to reduce the separate NTSC moiré problem.)

Another way to create a standard video signal from a native signal is with a standards converter, similar to those used for converting European television to American television and vice versa. The original image is stored in a large bank of RAM called a frame store, or buffer. Unlike a CRT, in which the image begins to decay when the electron beam passes on, a frame store holds each part of the image until it is replaced by a new frame. Meanwhile, the image is read out by a separate circuit operating at the rate of the converted signal, so the frame store acts as a buffer between the two video systems. No sync bars or other artifacts are visible except for minor discontinuities in movement.

A standards converter is the only effective way to get a standard video signal from a small Mac. But standards converters are expensive; buying a Mac II with a suitable video board would be cheaper than trying to convert a small Mac's video signal.

Getting a standard video signal out of a Mac II is fairly easy. You only need to add a genlock accessory video board to Apple's Macintosh II Video Card. Circuitry in the accessory board creates a standard video signal in two fields by extracting odd- and even-numbered lines successively from the frame store. When genlocked, Mac II images can be broadcast, recorded, and mixed with other video sources. The 480 lines produced by Apple's video card fit into the 484-line maximum of NTSC television.

A 480-line Mac II video image does, however, pose several problems when converted to NTSC format. The edges of the image (called overscan) will usually be hidden on a standard TV set. Because a single Mac frame is broken into two NTSC fields, fine horizontal detail will flicker. To reduce the flicker, objects that you put on the Mac screen should always have an even number of horizontal lines, and nothing should be less than two lines high. A progressive-scan television monitor can also help reduce flicker. Progressive-scan monitors always display 484 lines instead of alternating odd-numbered and even-numbered fields; they produce the "missing" lines either by repeating each line once (as in the Toshiba CZ2697) or by interpolating lines. Overall, line interpolation suppresses flicker better than line repetition.

New video boards now being developed for the Mac II by independent companies can digitize standard video images for processing by the computer and then convert the images back to standard form for editing into a television program. A Mac II could thus serve to generate video effects, at least for still frames; whether such a board can handle moving images will depend on its processing speed.

Mac II video boards that support large monitors generate more than 480 lines; viewing their images on standard video equipment requires standards conversion that averages or skips adjacent lines to reduce the line count so it falls within NTSC's 484-line maximum. Most of these boards generate more than 640 pixels horizontally, which a standard television monitor cannot resolve. Nonetheless, having additional pixels helps reduce jagged edges.

In the long run, higher-quality television could preserve all of the Mac's image quality. Some high-definition television systems can now meet or even exceed the Mac's requirements. Japan Broadcasting Corporation's experimental HDTV (high-definition television) uses 1125 interlaced lines (1041 visible lines) with a 20-MHz video bandwidth. To achieve such quality, however, HDTV requires the replacement of every component of present television systems, from cameras to recorders, transmitters to television sets.

If you need the best possible image of a Macintosh screen, either photographic or video, contact consultant John Monsour, 2062 Stanley Hills Drive, Hollywood, CA 90046.

23: Future Directions

Although the computer industry, especially the microcomputer industry, is changing fast, the short-term future is fairly easy to predict. Every manufacturer uses the same technologies; there are, after all, only a handful of important microprocessor chips and only a few ways to make a screen display. Everyone works with the same memory and disk-controller chips and the same disk drives. All the hardware components that will be commercially available in the next four to five years already exist today in prototype.

Within five years, every successful computer company will offer a full family of models, from powerful desktop designs to battery-operated portables. All the models will work together, using as similar an interface as the specific hardware allows. With little effort, you will be able to move files back and forth between any members of the family. Whenever you connect two computers, the operating system will automatically update files and make backups after checking the time and date of each file.

PROCESSORS

Someday, although you may not believe it now, you will need a faster computer; in a year or two, even the Mac II will seem slow. Making a computer run faster starts with faster central processing but ultimately demands overhauling the designs of disk drives, memory, and the bus. Developers can arrive at faster processing by a variety of paths—sometimes by several paths simultaneously.

Speeding up an existing CPU chip is the most straightforward route. Motorola's 68020 CPU chip in the Mac II, for example, can operate as fast as 16 MHz (it actually runs at 15.67 MHz), and versions that run at 25 MHz are already available. Software that runs on today's 68020 would need no modification to run on a faster version.

Improving the design of an existing chip so that it can process more complex instructions has produced Motorola's successor to the 68020, the 68030. This chip can move data more quickly to and from the bus, supports faster memory access, and has a larger data cache than the 68020 for information it is working on. The 68030 also incorporates the functions of the 68851 memory-management chip, without the delays resulting when a 68020 and a 68851 work together. Overall, the 68030 can process information about 75 percent faster than a 68020 for the same clock speed; its design should permit higher clock speeds as well. The 68030 can run all the same software as the 68020, although some programs will need modifying to take advantage of the 68030's greater processing power. A successor to the 68020 and 68030 will probably be announced in 1988 or 1989.

RISCs (reduced-instruction-set chips) take the opposite tack. Instead of processing more complex instructions, they process only a limited set of simple instructions—very fast. A complex maneuver that might consist of a single instruction on an ordinary CPU might take several RISC instructions to accomplish, but if the RISC runs fast enough, it can still execute the maneuver faster than its traditional competitor. Programs developed for traditional CPUs generally require considerable change to run on a RISC processor.

Several microcomputers are based on RISCs, including IBM's PC/RT and Sun Microsystem's Sparc computers (Mountain View, CA). Apple has also begun a RISC development project.

Adding specialized processors can take some of the computing load off the main CPU. With a few changes to take full advantage of each processor, existing software can usually support this design. For example, in the Mac II, a 68881 numeric coprocessor is already standard for speeding up long or complicated calculations. The newer 68882 numeric coprocessor works from 30 to 100 percent faster than the 68881, and an array processor—which is designed to process a block of information, such as a matrix, rather than a sequence of numbers—can speed calculations even further. Graphics coprocessors can greatly expedite processing for video displays and printers with few or no changes to existing software. (See Chapter 4.)

Parallel processing is a fundamentally different design that relies on a group of powerful general-purpose processors working together instead of a single central processing chip. An ordinary CPU chip can only process information sequentially, one step at a time; it is a processing bottleneck. A parallel-processing computer has more than one CPU—some have as many as 65,000—individually processing one step at a time but in aggregate processing many steps at a time. Full-scale parallel processing will probably not find its way into personal computers in the near future because the whole assemblage is very expensive, and existing software must be drastically modified.

A kind of parallel processor is already available for the Mac II, however. The TransLink board from Levco (San Diego, CA), which works as an accessory processor to the Mac II's 68020, contains up to four Inmos chips called Transputers (Colorado

Springs, CO). These chips operate in parallel: Adding more Transputers increases processing speed. Transputer programs must be specially written.

In 1988, users will be able to speed up their Mac IIs by installing a faster 68020 or a faster and more sophisticated 68030 chip—either on a speed-up board plugged into a NuBus slot or on the Mac II motherboard. Other fairly straightforward improvements in the drives and memory can make an improved Mac II run about three times faster than Apple's 1987 model. Some of these improvements are expensive, however, and, to the developer, not easily justified in a mass-market computer.

Paradoxically, faster processing won't necessarily mean that the computer's overall speed seems faster to the user. As processors become more capable, software will become more complex; a new computer might run four times as fast, but a complex new program might have it doing four times as much work.

MASS STORAGE

Floppy disks will remain widely used for years to come because the disks themselves are so inexpensive to make. Meanwhile, new technologies will enable storage densities to take a big jump—to five megabytes and more. Such densities will bring some problems, however. Most proposed 5 MB floppy drives can read but not write 800 KB and 1.6 MB formats. Dust will be a bigger nuisance than ever, and the mechanical alignment of the disk drives will need to be much more precise.

Hard disk drives will run faster and store more data; drives with hundreds of megabytes will be common. Virtually every desktop computer will have a hard disk; computers with floppies alone will be the exception.

Disk access speed in the Mac II can be improved with a better interface circuit. The SCSI chips in the Mac II support a maximum data transfer rate of 1.25 MB per second; new SCSI chips can support up to 5 MB per second. In Apple's original design, disk access is an input/output function that requires the CPU to feed information to the disk relatively slowly. A direct memory addressing (DMA) interface could set up a section of memory to handle disk transfers. The CPU merely needs to write a block of information to that memory location and can then do other work; the disk controller grabs the information and transfers it to disk.

Developments in optical discs and digital audio tape are discussed in Chapter 6.

DISPLAYS

The long-established CRT will remain dominant in desktop computers through this decade. No other display technology has equaled its relatively high contrast and resolution and its low manufacturing cost. Only in portable computers will flat-panel technologies dominate—and then only because size, weight, and power constraints are more important for portables than pixel count or price.

Large displays will get larger and gain pixels. The practical limit for a CRT, whether monochrome or color, is likely to be about 3000 by 2000 pixels. Screen resolution will probably stay at about 70 to 150 dpi. Although 5000-by-3000-pixel CRTs with 300 dpi and more have been built for specialized applications such as typesetting and creating images for microfilm reproduction, costs have so far been prohibitive. The resolution of flat-screen displays will not be any better, so computer screens will continue to be low-resolution devices for the foreseeable future.

NETWORKS

Networks will become common and interlinked by the end of the decade. Present local area networks will include gateways to telephone systems and, thus, to other networks. Electronic mail and file sharing, whether with the next office or another continent, will be embedded in computers as part of the operating system. You won't even have to keep track of where a particular file is located; the system will automatically find the most recent version.

TELECOMMUNICATIONS

ISDN (Integrated Services Digital Network) has begun to replace today's analog telephones. The changeover to ISDN will take until well into the twenty-first century to complete, however, for it requires a massive investment in optical fiber links among telephone exchanges. Fortunately, few local loops—the wires between homes and offices and the local exchange—will need replacement. ISDN will at last offer decent communication speed—64 kilobits per second, instead of today's 20 kilobits per second under the best conditions. A double-spaced typewritten page will transfer in ⅓ second; even the most complex small Macintosh screen will transmit in 3 seconds, and the majority of screens will take much less time.

ISDN supports multiple voice and data channels on the same phone line, using a coding device at the telephone and a complementary decoder at the local exchange. Multiple data channels can be useful for burglar alarms and other sensing devices. The phone companies are delighted; they can charge you for several calls simultaneously—your own and those made automatically by the sensors. (For more technical information about ISDN, see Chapter 20.)

Faster data communication over phone lines will make graphical interfaces for dial-up information services practical, so these services will finally become easier to use. They will then attract more customers, in turn bringing the cost low enough to attract still more.

Everyone also hopes that the advent of new high-speed links will help standardize communication protocols, but this is unlikely. New software will at least insulate individual users from the protocol tangle by automatically detecting which protocol is required and adjusting as needed.

GETTING INFORMATION INTO COMPUTERS

In this decade, the keyboard, with help from a mouse, will remain the predominant method of entering information into a computer. Everyone wants a computer that will recognize human speech, but the challenge is formidable.

With present technology, modestly priced units can recognize a limited vocabulary of isolated spoken words, but these units are speaker-dependent—they must be "trained" to recognize individual speakers. Expensive systems can interpret limited connected (but not continuous) speech with a vocabulary of up to 500 words.

A powerful system able to understand standard spoken business English will probably not be available until well into the 1990s. Such a system will require a sophisticated program to reduce the many ambiguities inherent in spoken English.

SOFTWARE

Although future software will still do the tasks you expect software to do today, how it does them will change.

Live links, discussed in Chapter 16, will become increasingly common. With live links, information created in one application and moved to another is updated automatically. If you create numbers with a spreadsheet program and then move them into a word-processing document, any changes you make in the spreadsheet will be automatically carried to the document. Eventually, all software will be linked.

A second major direction will be toward natural language—ordinary English. To retrieve information from a database, you will no longer have to learn a rigid series of commands. Instead, you'll be able to type a question in plain English; the software will analyze the syntax of your question (parse it) and figure out what information you want to know. Parsing is akin to sentence diagramming combined with a dictionary of meanings. With parsing, computers can for the first time be said to "understand" language.

Microcomputer databases that accept simple natural-language input began appearing in 1984 (for the IBM PC), but that was just the beginning. The twin steps of parsing and linking will open the way to a new era in software.

DATABASE MACHINES

Everyone works by building upon old data. But compared with computers, people are disorganized; computers need organized data. With today's computers, you spend an inordinate amount of time organizing information so that computers can understand it.

Letters, for example, begin with a name and address that are already stored in an address book or file. Even if you are using present software to store addresses and write letters, you must still explicitly search for an address or an old letter before writing a new one. Future software will search automatically; it will parse your input as you type.

If you type *John Doe*, the software will immediately put John Doe's address and telephone number in a secondary window. If there is no John Doe in your address file, the program will wait until you type in a new name and address.

Regardless of where or how you type the address, the software will detect an address by scanning for terms such as *Ave.*, state abbreviations, and so forth, and then look in nearby text for further address information. The information will be automatically organized into address format; you will have a chance to check it, and it will then be added to your name-and-address database.

If your letter begins, "In response to your letter of last week . . . ," the software will automatically figure out that "last week" means April 8 through 14 and will check for letters from Doe. (All letters will arrive via electronic mail, of course.) If no such letter is found, the software will expand its search, looking further back in time. If you regularly underestimate the date, the program will automatically start looking earlier.

If you refer to a company's name in the letter, all the information in your files about that company will be available instantly. The software will anticipate all common data needs, accepting natural-language guidance for less common requests ("What is the telephone number of the hotel closest to the meeting?").

Software will be unstructured. You will not have to ask for a word processor or a database when you start working; one large, seamless program will analyze and suitably organize your work. All program functions will always be available without arbitrary boundaries between them. The software will analyze and adapt to your specific working habits. If you enter calendar dates in a characteristic way, the program will adopt the same format, but it will retain the ability to read dates from other sources in any format. On-line dictionaries will always check spelling and offer corrections for common mistakes.

Software like this will be extremely complex. Unlike some star programs of the past, which were written by lone programmers working nights, this new software can be written only by large, well-organized programming teams.

FINDING OUT ABOUT NEW PRODUCTS

You will have no difficulty finding out about new products for the Mac; the problem will be sorting out the good from the bad. The microcomputer industry is putting out a deluge of announcements touting new Macintosh accessories and software. Claims and counterclaims crowd magazine advertisements.

Pre-announcements

Many companies announce products early—months, sometimes a year or more before they actually deliver. They may want to preempt the field and discourage people from buying competing products, or they may simply be optimistic; even competent companies can run into last-minute snags.

Magazines encourage such pre-announcements by trying to be the first to have the latest news. For major products, a magazine may strike an unspoken deal with the company, trading a splashy article or cover in exchange for early information. Few magazines publish corrections when the hot news turns out to be mostly hype.

Computer engineers tell many funny stories about last-minute crash projects done to simulate a product for the press or investment analysts. The demonstration may have nothing to do with the product—usually it's a lot of fancy screen graphics—but it almost always convinces the gullible.

Pre-announcements are often the result of scheduling problems. Magazines typically plan articles and sell advertising space for each issue three to four months ahead of time. Many products will fail if they do not start selling successfully within a few months of the date they are released; if sales are delayed, competing products will appear. So, to make the publicity coincide with the product's release date, some companies start the announcement process months before the product itself is ready.

The best companies take a more conservative approach, announcing products only when they are ready to ship. All other things being equal, these companies deserve more support than those that regularly jump the gun.

Rumor mills also generate many stories about forthcoming products. Many rumors are inaccurate, based on wishful thinking. When a company plans a new computer, it will usually design several alternative configurations; most will never be built for sale, but they will generate additional rumors. Sometimes the rumor mills criticize an unfinished, unannounced product for not working well; such criticism is unfair, for if the product were working well, it would be finished.

Problems with New Products

Should you always look for the latest and most glamorous product? No, especially if you depend on the Macintosh in your work. Instead, concentrate on products and companies that have an established reputation. Ask other people what their experience with a particular product has been. You might not always be the first with something new, but you won't be an unwitting guinea pig either.

New software often does not work quite as well as the developer planned. Bugs, or errors, are common, particularly in early versions. You can learn to live with some bugs, especially if they afflict only little-used features. More serious are major bugs that affect a program's utility. A database program with a delete-record command that doesn't work, for example, can only lead to frustration and unpredictable results. Such major bugs aren't unusual; several bestselling programs for other computers have had such faults for years. Sometimes the seller will try to explain away the bug as an "undocumented feature."

The chance of bugs increases with a program's complexity. Often, very complex programs may never be completely debugged; for the developer, it becomes a question of how severe the bugs are and whether it is worth the effort to fix them. Frequently, the bugs are fixed only in a new version that adds features—and new bugs. Nevertheless, the most highly developed programs should be bug free for all practical applications.

MICROS AND MAINFRAMES

Are microcomputers really taking over more and more computer tasks? Are they replacing minicomputers and mainframes? Can they replace minis and mainframes? How does a microcomputer compare with a mainframe in power?

The design of mainframes and minicomputers differs too much from that of micros to make simple comparisons possible. Minicomputers and mainframes must usually accommodate multiple users simultaneously and therefore have far more sophisticated input-output arrangements than micros. The most common measure for computing power—millions of instructions per second (MIPS)—is an interesting, but by no means comprehensive, criterion; it tends to downplay overall differences in throughput. What's more, some companies exaggerate how many MIPS their computers can process.

Among current computers, a Macintosh SE handles about 0.5 MIPS; a Mac II handles 2.4 MIPS. A DEC VAX does 1 to 6 MIPS, the larger IBM mainframes around 15 to 30 MIPS, and the most advanced supercomputers reach 150 MIPS. By 1990, large computers will improve by 3 to 10 times. Micros will grow faster: An enhanced Mac II with a 25 MHz 68020 central processor could do 4 MIPS; with a 68030 processor, it could reach 5 or 6 MIPS. By 1990, microcomputers could do 15 to 25 MIPS, particularly if they incorporate auxiliary processors.

Why are microcomputers developing faster than their larger cousins? Developing a microcomputer takes much less time than developing a mainframe. Designing the central processing chip consumes the most time; once the chip is available, a complete microcomputer based around it takes only about a year to design and build. In contrast, developing a new mainframe takes five to six years. The major mechanical problem in mainframe development has become how to get rid of the heat generated by the circuitry.

Micros have modest computing power but are easy to use; mainframe computers have lots of raw computing power but are extremely difficult to use. Ultimately, micros are likely to have about one-tenth the computing power of mainframes and a far superior cost-performance ratio.

In 1988, about 90,000 mainframe computers were operating worldwide—roughly a week's production of microcomputers. In the long run, mainframes will become utility devices for micros, serving for complex programs and large database needs. Users will see only the microcomputer interface, which will hide any difficulties associated with mainframe operation.

Appendix: Technical Topics

This appendix is for readers with some technical expertise and knowledge of computers. Some of these topics have been brought up in the body of the book, and more detail is presented here; some topics are supplements. Other technical references include the *Technical Introduction to the Macintosh Family* and *Macintosh Family Hardware Reference,* two books prepared by Apple Computer and published by Addison-Wesley (Reading, MA).

STANDARD VIDEO SCREEN

The video screens on all the small Macintoshes have the same characteristics:

Video screen: 9-inch diagonal, black-and-white.

Pixels: 512 by 342 pixels, noninterlaced.

Image size: 4¾ by 7 inches (12 by 17 cm).

Phosphor: P4, medium-fast, dual component.

Spectrum peaks at 460 and 560 nanometers.

1931 CIE color coordinates: $x = 0.27$, $y = 0.30$.
1964 CIE color coordinates: $u = 0.18$, $v = 0.29$.
Equivalent color temperature: 11,000 degrees Kelvin.

Horizontal scanning frequency: 22.254545 kHz.

Horizontal flyback time: 13 microseconds.

Frame rate: 60.1474 Hz.

Vertical blanking interval: 1.24 milliseconds, 7.5 percent of duty cycle.

(continued)

continued

Dot rate (frequency of generating dots on screen): 15.6672 MHz.

Scanning time per pixel: 65 nanoseconds.

Frame rate: 60.15 Hz (the speed that falls out from dividing the master clock frequency).

Video bandwidth: 20 MHz. This is the bandwidth necessary for an external video amplifier; no internal video amplifier is present.

There are too many different Mac II video screens to list specifications here.

MICROFLOPPY DRIVES AND DISKS

Three different types of 3½-inch microfloppy disk drives exist: the original version developed by Sony of Japan, a version that emulates 5¼-inch drives, and the type made by Sony for Macintosh and Apple II computers. All come in single-sided and double-sided models, and all but the original Sony design come in standard and high-density models as well. The drive installed in Apple computers has a variable-speed disk controller and an autoeject mechanism; it is the only one that can be used with a Macintosh.

Apple records data on floppy disks in a unique way, using a technique originally developed for the Apple II and implemented by a chip called the Integrated Wozniak Machine. The company is now developing a new controller circuit that can handle both its own and standard recording methods, which will allow high-density microfloppy drives in the Mac to read and write both Mac and MS-DOS or OS/2 disks.

The heads of double-sided disk drives are offset; the top one is slightly closer to the center (by 59 mils, or 1.5 mm: the equivalent of eight tracks). On early larger floppy disk drives, the offset was supposed to improve head contact with the disk, but new head geometries probably make it unnecessary.

Microfloppy disk specifications are as follows:

Disk jacket size: 3.54 by 3.7 by 0.13 inches (90 by 94 by 3.3 mm).

Disk jacket construction: ABS plastic (Sony brand).

Disk inner diameter: 25 mm.

Disk outer diameter: 86 mm.

Recorded area (bottom only on single-sided disks):
- Bottom: inner radius, 22.428 mm; outer radius, 41.961 mm.
- Top: inner radius, 21.666 mm; outer radius, 40.970 mm.
- Tracks: 80 tracks per side; track density, 135 tracks per inch.
- Track width: 0.12 mm.
- Substrate: polyethylene terephthalate (Sony brand).
- Coating: cobalt adsorbed iron oxide, 0.0019 mm thick (Sony brand).

Error rate:
- Soft error rate (correctable on retry): 1 in 10^9 bits.
- Hard error rate (not correctable): 1 in 10^{12} bits after 10 retries.
- Seek error rate (finding the right track; correctable with retry): 1 in 10^6 seeks.

DISK DRIVE SPEED

This section supplements the discussion in Chapter 6.

- Head motor. One of the critical factors governing the access time of a hard disk drive is the type of motor that moves the head assembly. Slower hard disk drives (average access times of 65 ms or longer) move the heads by means of stepper motors. Medium-performance disk drives (average access times of 30–45 ms) use linear motors. Expensive, higher-performance drives (average access times less than 30 ms) move the heads as if they were the cones in a loudspeaker, by means of an electromagnet called a voice coil. Servo signals recorded in the narrow tracks help the heads find their destination.
- Track geometry and cylinders. For a given number of kilobytes, it takes a double-sided disk drive half the number of head movements—and thus less time—to write on a double-sided floppy disk as on a single-sided disk because the two recording heads move together, writing contiguous tracks alternately on the bottom and top surfaces. Looking at it another way, a double-sided drive can store more per head movement on a double-sided disk than on a single-sided disk. In the middle tracks, for example, a drive can store 5.12 KB on a single-sided disk before it must step to the next track; on a double-sided disk, it can store 10.24 KB. (This alternating bottom-and-top writing pattern explains why a single-sided drive cannot read useful information from a double-sided disk. In contrast, a double-sided drive can simply shut off its upper head to format, read, and write single-sided disks.)

  Hard disk drives take the same idea further: An array of heads moves together to write on a stack of magnetic platters sharing a common spindle. All the platters are magnetically coated on both sides, and there is one head for each platter surface. When recording a file, the drive steps the multiple-head assembly to a specific track location and then writes in the tracks at that location on both sides of each platter in turn before moving the heads to the next track. The tracks on all the platters associated with a particular head position are collectively known as a cylinder.

  In contrast with a floppy disk drive, which can store 5.12 or 10.24 KB, a hard disk can typically store from 35 to 100 KB before stepping to the next track. To double a cylinder's capacity, some expensive hard disk drives have two heads for each platter surface. This also reduces the disk drive's average access time because the total distance a head has to travel is cut in half; one head covers the tracks on the inner half of a platter, and the other head covers the outer tracks.
- Block size and fragmentation. A disk operating system allocates the storage available on a disk into logical units called blocks (distinct from sectors, tracks, and cylinders, which are physical units), and the file allocation table in the disk directory keeps track of which files are in which blocks. The smallest

practical block is usually equivalent to a sector; the largest blocks can be bigger than a cylinder. A block is the smallest space that a file can occupy on a disk. Small files are padded with null characters to fill out a block; large files take up many blocks.

On every floppy disk, blocks 0 and 1 are reserved as boot blocks; the disk directory is in block 2. All other blocks are assigned independently. Hard disk drives are set up similarly but may allow more space for the directory. On a fresh disk, a large file will occupy contiguous blocks, but after many files have been created and deleted, files tend to fragment into noncontiguous blocks as the operating system stuffs data anywhere it can. If block size is small, fragmentation gets worse, and more head movements are required to reach each block of a given file. If block size is large, the disk cannot store many small files because the blocks are quickly used up and store more padding than data.

- Latency. Once a head has reached the correct track, it has to wait for the desired information to spin around; this delay is called latency.
- Blind writes. Information sent to a hard disk drive connected to a SCSI port can be checked during the writing process to make sure it is accurate. But checking slows the writing process. In a blind write, the information is not checked, and writing is speeded up. According to Apple, blind writes are more likely to cause trouble on a Mac Plus than on a Mac SE or Mac II. A Mac Plus can transfer data through its SCSI port at 312 KB per second in a blind write and at 142 KB per second with checking.
- Moving a hard disk from one Mac to another. If you move a hard disk from a Mac II (1:1 interleave) to an SE (2:1 interleave), it will perform poorly because the SE controller cannot keep up with a 1:1 interleave and so has to let the disk spin one complete revolution for every sector it reads. To correct this problem, reinitialize the disk and reinstall driver information for an SE—but back up your files first.

INPUT/OUTPUT PORTS

For readers who want to wire their own cables for connecting peripherals and other computers to the Macintosh ports, the following sections offer directions.

Serial Ports

The Mac's RS-422 serial ports are specified for operating as fast as 230,400 bps. With external clocking, the ports can run as fast as 920,000 bps. The RS-422 interface uses balanced signal lines in pairs; the related RS-423 interface uses unbalanced lines.

Although the serial ports in all Macs operate in much the same way, their features and plugs differ. The 128 KB Mac and the Mac 512K and 512KE had a 9-pin D-connector;

the Mac Plus and later models have an 8-pin mini-DIN. Table A-1 summarizes the differences in serial ports and wiring for all Mac models.

Table A-1. MACINTOSH SERIAL PORT SPECIFICATIONS

	Mac SE, Mac II	*Mac Plus*	*128 KB Mac, Mac 512K, 512KE*
Plug type	Mini-DIN	Mini-DIN	9-pin D
Data carrier detect (DCD)	Pin 7	Not available	Not available
Clock input	Pin 7 (modem port only)	Not available	Not available
Handshake out (DTR, data terminal ready)	Pin 1	Pin 1	Not available
Handshake in (CTS, clear to send)	Pin 2	Pin 2	Pin 7
RxD− (receive data)	Pin 5	Pin 5	Pin 9
RxD+ (receive data/balanced)	Pin 8	Pin 8	Pin 8
TxD− (transmit data)	Pin 3	Pin 3	Pin 5
TxD+ (transmit data/balanced)	Pin 6	Pin 6	Pin 4
Signal ground	Pin 4	Pin 4	Pin 3
Chassis ground	Shell	Shell	Pin 1 and shell
+5V power	Not available	Not available	Pin 2
+12V power	Not available	Not available	Pin 6

On a Mac SE or a Mac II, pin 7 serves two purposes. It can accept a data carrier detect signal, usually sent by your modem to say that a remote modem needs control of a phone line. On the modem port, pin 7 can also accept an external timing signal from a synchronous modem. The computer gives the modem port higher priority than the printer port, so it is more suitable for some medium-speed and high-speed modems. (Low-speed modems—2400 bps or slower—can be connected to either port.)

RS-422 protocols define the lines for transmitting and receiving data in balanced pairs. For short cables and for connections to RS-232C or RS-423 devices, only one side of the balanced pair need be connected (RxD– and TxD–). For long cable runs between RS-422 devices, the balanced pair is more resistant to electrical noise and interference.

Mini-DIN connectors are hard to find and harder to solder than 9-pin D-connectors. If you need an adapter from one to the other, here is how to wire one. The lines indicate connections between the pins listed by number in each column.

Mini-DIN		*9-pin D*
2	———	7
3	———	5
4	———	3
5	———	9
6	———	4
8	———	8

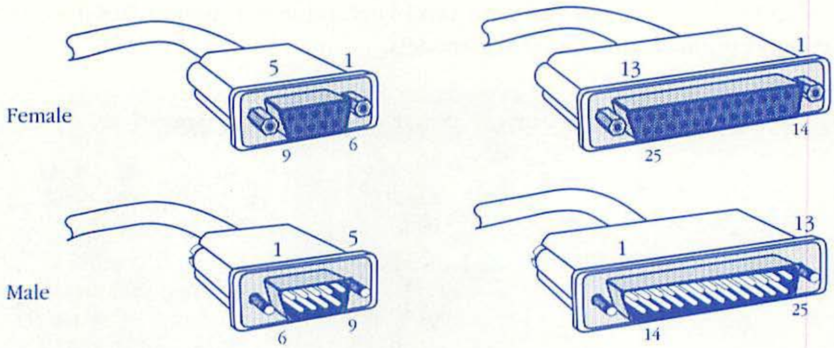

Wiring for 9-pin and 25-pin D-connectors, viewed from the contact side.

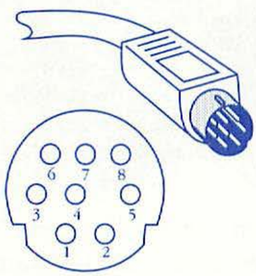

Mini-DIN cable connector (male).

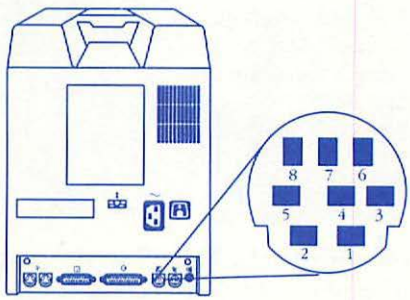

Mini-DIN chassis connector (female).

To connect one Macintosh to another, use a null-modem cable wired according to the scheme below. Pins are listed by number and their assignments are given on the left; NC means no connection.

	Macintosh A		*Macintosh B*	
	Mini-DIN	*9-pin D*	*Mini-DIN*	*9-pin D*
Ground*	4	3	4	3
Transmit data	6	4	8	8
Transmit data*	3	5	5	9
Receive data	8	8	6	4
Receive data*	5	9	3	5
Handshake out	1	NC	2	NC
Handshake in	2	NC	1	NC

For short cable runs with software handshake, you need to connect only the three wires marked with asterisks.

The following schemes show the wiring for cables between a Macintosh (left two columns) and some common computers and peripherals (right column). The connections are labeled on the left from the Mac's point of view: "Transmit data" means that the Mac is transmitting and the other device is receiving.

For connecting a Macintosh to an ImageWriter or Tandy Model 100 portable computer:

	Mac		*ImageWriter*
	Mini-DIN	*9-pin D*	*25-pin D*
Ground/balanced line	4 & 8	3 & 8	7
Receive data	3	5	3
Transmit data	5	9	2
Handshake	2	7	20

For connecting a Macintosh to a Hayes modem:

	Mac		*Hayes modem*
	Mini-DIN	*9-pin D*	*25-pin D*
Signal ground	4	3	7
Transmit data	3	5	2
Handshake	2	7	5
Receive data	5	9	3
Chassis ground	Shell	1	1

For a null-modem cable to an IBM PC or PC/XT with a 25-pin D-connector, wired DTE:

	Mac		*IBM PC*
	Mini-DIN	*9-pin D*	*25-pin D*
Ground	4	3	7
Transmit data	3	5	3
Receive data	5	9	2
Handshake	1	NC	6

On the IBM PC, pins 6, 8, and 20 should be jumpered together, and pins 4 and 5 should also be jumpered together. The handshake line is unnecessary if handshaking is done by the communication programs (XON/XOFF).

For a null-modem cable to an IBM PC/AT and PS/2 with a 9-pin D-connector, wired DTE:

	Mac		*IBM PC/AT, PS/2*
	Mini-DIN	*9-pin D*	*25-pin D*
Ground	4	3	5
Transmit data	3	5	2
Receive data	5	9	3
Handshake	1	NC	6

On the IBM PC/AT, pins 6 and 1 should be jumpered together, and pins 7 and 8 should also be jumpered together. Again, the handshake line is unnecessary if handshaking is done by the communication programs.

Apple DeskTop Bus

The pin assignments for the Apple DeskTop Bus (Mac SE and Mac II) are:

1	Data
2	Power-on (Mac II only)
3	Power +5 volts; all connected ADB devices should draw no more than 500 mA
4	Signal ground

Shorting pin 2 to ground (pin 4) momentarily—what the Power-on key on the keyboard does—turns on the Mac II. You can turn on a Mac II remotely if you have an auto-answer modem that closes a relay on answering a call; connect the relay across pins 2 and 4.

The Apple DeskTop Bus connector is the same as the Y/C connector used by Super VHS and ED-Beta videocassette recorders.

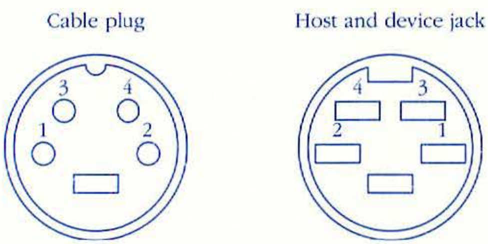

Keyboard Connector

The pin assignments for the keyboard connector on the Mac Plus and earlier models are:

1	Black	Ground
2	Red	Clock
3	Green	Data (device ID and keyboard scan code)
4	Yellow	+5 volts

The keyboard cable uses the same plugs as a modular telephone handset. The cables are not interchangeable, however. The keyboard cable is heavier gauge and wired straight through (pin 1 to pin 1, and so forth), whereas the handset cable wiring is flopped (pin 1 to pin 4, pin 2 to pin 3, and so forth).

Mouse Connector

The pin assignments for the mouse connector on the Mac Plus and earlier models are:

1	Ground
2	+5 volts
3	Ground
4	X-2
5	X-1

(continued)

continued

6	Not connected
7	Switch bar for mouse button; pressing the button connects this line to ground
8	Y-2
9	Y-1

The X and Y connectors are the quadrature signals for the horizontal and vertical axes. If, for some reason, you want to reverse the orientation of the mouse's movements (reversing left-right or up-down movements), you can swap X-1 and X-2 or Y-1 and Y-2.

Disk Drive Port Connector

The wiring and operation of the floppy disk drive port is the same on all Macintosh models:

1–4	Ground
5	−12 volts
6	+5 volts
7 & 8	+12 volts
9	Not connected
10	Motor speed control
11	CA0 (status control line)
12	CA1 (status control line)
13	CA2 (status control line)
14	LSTRB (status control line)
15	Write request
16	SEL (select disk drive)
17	External drive enable
18	Read data
19	Write data

NuBus Connector

The 96 pins in a NuBus connector are distributed as follows:

4 pins for utility purposes such as resetting the bus.

4 pins for control.

32 pins for bus address and data.

5 pins for bus arbitration.

2 pins for parity checking (not used in a Mac II).

4 pins for slot identification.

23 pins for power.

11 pins with +5 V; 12 A maximum for all slots.
2 pins with +12 V; 2.5 A maximum for all slots.
2 pins with −12 V; 1 A maximum for all slots.
8 unused power pins. (In the original NuBus specification, these pins would have carried −5.2 V.)

22 pins for ground. (Having multiple ground wires reduces electrical noise and interference.)

Mac SE Bus Connector

The single expansion slot in a Mac SE uses a 96-pin connector following DIN specification 41612 and MIL-C-55302. The CannonITT connector type is G06.

The SE Bus differs from NuBus in many ways. The bus address is carried on 23 pins that are separate from the 16 pins that carry bus data. Power is available as +5 V (1.5 A max), −5 V (100 mA max), +12 V (150 mA max), and −12 V (100 mA max).

SOUND GENERATION

For a small Mac, the sound generator operates during what is called horizontal flyback time. In other words, the 68000 CPU processes sound-generating information when it is not busy with the screen—while the electron beam is traveling back ("flying back") from the right edge of the screen to the left edge. Flyback occurs at the horizontal sweep rate, or every 44.93 microseconds, so all sound frequencies are multiples of this time. The highest possible frequency is twice this period, 89.96 microseconds, or 11.116 kHz (Nyquist limit).

From the programmer's point of view, there are three different sound synthesizers:

- The four-tone synthesizer for harmonic tones; 8-bit digital-to-analog conversion.
- The square-wave synthesizer for beeps and other effects.
- The free-form synthesizer for complex music and speech.

On a Mac II, sound is generated by a separate sound chip and is subject to fewer limitations.

BATTERY

All Macs have a battery to power a built-in clock/calendar and parameter RAM. These circuits remain alive at all times, even when the computer itself is turned off. In models before the Mac SE, the user can change the battery (Eveready No. 523 or equivalent). In the Mac SE and II, the lithium battery soldered into a circuit board is supposed to last for seven years; changing this battery normally requires a technician.

In a Mac II, the same battery also supplies current for an internal power-on switch that starts the computer when you press the Power-on key on the keyboard. If the battery fails, the Mac II will not turn on.

MEMORY

This section gives more details about the contents of parameter RAM and the possible configurations of main RAM.

Parameter RAM

Parameter RAM contains 256 bytes. In addition to time and date, it stores the following information:

Validity status: a check that the last information written into parameter RAM was valid; if not, default values are restored.

Modem port configuration (Port A): Default is 9600 baud, 8 data bits, 2 stop bits, no parity.

Printer port configuration (Port B): same as for modem port.

Time setting: Default is midnight, January 1, 1904; the clock can keep time until the year 2040.

Default application font: font 2.

Keyboard repeat threshold: Default is 24 ticks, settings in 4-tick steps (one tick is $1/60$ second).

Keyboard repeat rate: Default is 6 ticks, settings in 2-tick steps.

Printer port: whether printer information is sent to the printer or the modem port; default is printer port.

Volume control: for speaker.

Mouse double-click time: Default is 32 ticks, three settings in 4-tick steps.

Mouse scaling.

Pointer blink time: 32 ticks, settings in 4-tick steps.

Menu blink: how many times a menu selection will blink (0 to 3).

Start-up disk drive: internal or external floppy disk; on an SE or II, also stores the SCSI address of the start-up disk.

These values are read from parameter RAM into main RAM when you turn on the power. The Control Panel desk accessory—and the public-domain program PRAM, available from user groups—allows you to change most of these values.

Occasionally you may need to erase the contents of parameter RAM to restore the default settings. On a Mac Plus or earlier model, you simply remove the battery and then replace it after a few minutes (a capacitor holds power for a few minutes even after the battery is removed; the delay lets the charge dissipate). To reset the parameter RAM of a Mac SE or a Mac II, press the Option, Command, and Shift keys while opening the Control Panel; you will lose all the information in parameter RAM except the time.

Memory Configurations

A Macintosh Plus or SE has four SIMM (single in-line memory module) sockets for memory chips. SIMMs must be installed a pair at a time in row 1 (sockets 1 and 2) and row 2 (sockets 3 and 4). The following configurations are possible:

Total memory	*Memory chip size*		*Resistor A*	*Resistor B*
	Row 1	*Row 2*		
512 KB	256 Kb	None	150 ohms	150 ohms
1 MB	256 Kb	256 Kb	150 ohms	None
2 MB	1 Mb	None	None	150 ohms
2.5 MB	1 Mb	256 Kb	None	None
4 MB	1 Mb	1 Mb	None	None

A configuration with only 512 KB total memory—less memory than the computers are usually shipped with—is uncommon.

The two resistors signal row and memory chip size to the controlling circuitry. Resistor A corresponds to R35 on a Plus, R8 on an SE; resistor B is R36 on a Plus, R9 on an SE. The memory chips must have an access time of 150 nanoseconds (ns) or faster; all the chips in a row must be the same speed. Faster memory chips will not improve the speed of the computer.

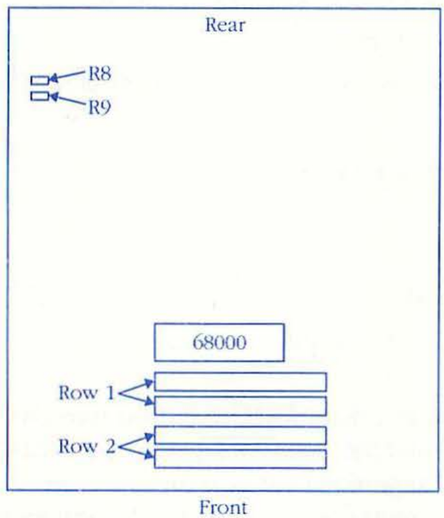

SIMM locations in a Mac Plus.

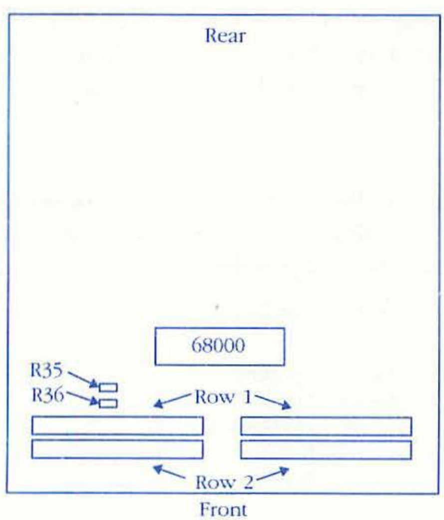

SIMM locations in a Mac SE.

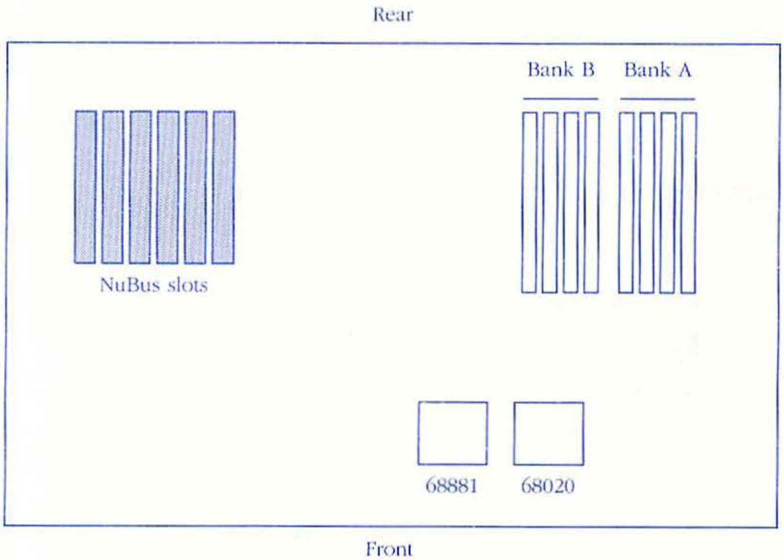

SIMM locations in a Mac II.

Mac IIs can accept up to eight SIMMs. They are installed in two banks, A and B, of four SIMM sockets each. Bank B is closer to the NuBus slots. All memory must be installed a bank at a time in the following configurations:

Total memory	*Bank A*	*Bank B*
1 MB	256 KB	None
2 MB	256 KB	256 KB
4 MB	1 MB	None
5 MB	1 MB	256 KB
8 MB	1 MB	1 MB

The Mac II automatically checks for the presence and capacity of memory; it does not need resistors. The memory chips must have an access time of 120 ns or faster. Memory can also be installed in NuBus cards. The first Mac IIs had a bug in ROM that prevented the computers from recognizing more than one megabyte of RAM on a NuBus card. The problem was solved in later production, and Apple offers a replacement ROM for owners of the first machines.

THE LASERWRITERS

The LaserWriter RS-422 port is designed to be used only with AppleTalk, not for direct connection. The direct-connect RS-232 port is wired DTE as follows:

2	Transmit data
3	Receive data
4	Request to send (optional, only if needed by host)
7	Signal ground
20	Data terminal ready (optional, only if needed by host)

When a LaserWriter starts up, it prints a page that gives its status (you can suppress this page with a short PostScript program available from user groups).

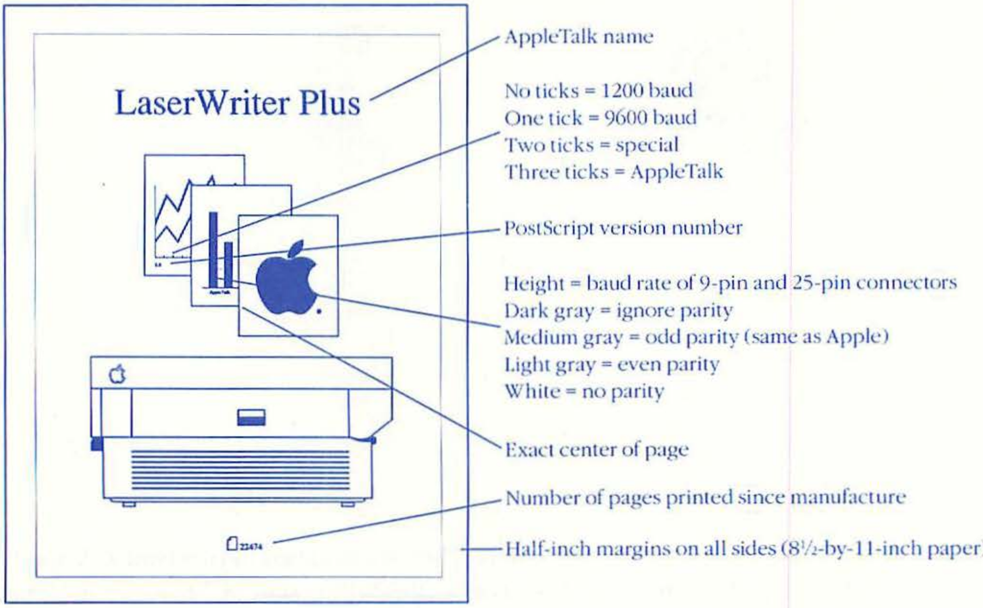

SYSTEM ERROR CODES

You will sometimes see an alert box containing a bomb icon. This box contains an error ID number that is generally useful only to the programmer, but if you report a problem to the software developer, you should supply this ID number. For the curious, here are the error types:

1 Bus error: can occur only on a Mac II or Mac XL.
2 Address error: Program referred to an odd address (usually trying to address a word on a byte boundary).
3 Illegal instruction: The 68000 did not recognize an instruction.
4 Division by zero.
5 Range check error: numerical range checking.
6 Overflow: integer overflow.
7 Privilege violation: refers to 68000 modes not used in the Macintosh.
8 Trace mode error: for debugging programs.
9 Line 1010 trap: unlikely but serious system error.
10 Line 1111 trap: for debugging programs.
11 Hardware exception error.
12 Unimplemented core routine: undefined system code.
13 Uninstalled interrupt: Interrupts are signals that grab the CPU's attention when it is doing something else.
14 I/O core: report of a very low-level I/O error.

(continued)

continued

15	Segment loader error: failure of an attempt to read a program segment into memory from disk.
16	Floating point error: a numerical computation problem.
17–24	Can't load package: missing segments of the system file such as disk initializer, minifinder, and so on.
25	Out of memory.
26	Bad program launch: could not find a program.
27	File system map trashed: faulty information about the files on a disk.
28	Stack ran into heap: The stack, which passes parameters and allocates variables, has run into the heap, which holds the application program.
30	Disk insertion error.
31	No disk insertion.
32–53	Memory manager error.
41	The file named Finder can't be found on the disk.
32767	General system error.

DISASSEMBLY

According to Apple, only its dealers and representatives are supposed to open a small Mac. They can install memory, hard disk drives, and accessory boards in the SE without voiding Apple's warranty as long as the computer is not physically modified. If anyone else opens the case, the warranty may be voided. You have to decide. If the work is done with care, there is no way to tell if an "unauthorized" person has opened the case. The warranty is good for only 90 days, anyway.

Whatever you think of Apple's recommendations, you should open a Macintosh case only for good reason and if you have the necessary technical skills. The inside of a small Mac is more hazardous than an Apple II or IBM PC or PC clone, for it contains the high-voltage power supply that drives the video display and the fragile CRT. You can get a dangerous shock even if the power cord is unplugged. Do not attempt disassembly unless you have electronics experience and know how to work inside a television set.

To disassemble a small Mac, disconnect all cables and remove the INTERRUPT/RESET button if it is installed. Put the Mac face down on a soft cloth and remove four screws from the back (use a Torx T-15 screwdriver with a 6-inch-long shaft). Small Mac models except the SE have a fifth screw underneath the battery cover. Pulling off the back cover requires some effort; two people can make the job easier. Be especially careful while the back is coming off because the rear of the CRT is not supported.

So what can you do after you've opened it?

For the early Macs, there isn't much you can do unless you have the parts and skills necessary to perform repairs or upgrades. The same considerations apply to a Mac Plus, except that adding memory is simple. For a Mac SE, you can add memory, install an accessory board, or replace the fan with a quieter one.

Installing and removing SIMMs is not difficult; they slide in and out of two locking plastic guides. If you are uncertain about the procedure, have someone show you how to open and close the guides.

If your SE has a noisy fan, you can replace it with an SE Silencer fan from MOBIUS (Oakland, CA). The Mac SEs produced in 1987 were particularly noisy, but by 1988 Apple was building them with quieter fans. Macs before the SE did not have a built-in fan; if you add equipment inside these models, you may want to add a fan as well to cool the circuitry.

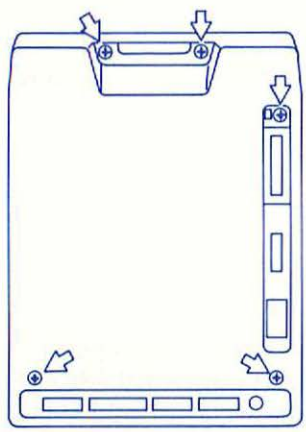

All small Mac cases are held together by four screws, two each at the top and bottom; the Mac Plus and earlier models have a fifth screw behind the battery compartment door.

The analog circuit board along the left side of a small Mac contains several controls and a fuse. Again, don't attempt any adjustments unless you know exactly what you are doing.

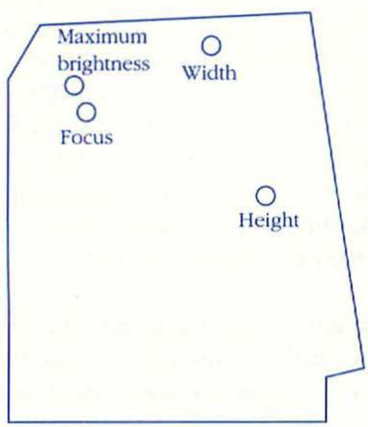

Mac Plus and earlier models.

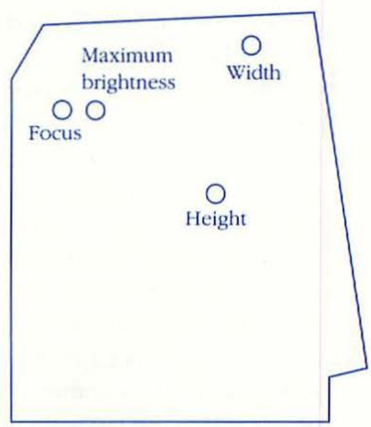

Mac SE.

Video adjustments for small Macs: This view is from the foil side. To adjust, remove the insulating sheet and insert a nonconductive screwdriver through the holes in the circuit board.

Reassembly is straightforward. Gently tighten the two screws under the carrying handle; they go into plastic. The other screws go into metal threads.

A Mac II case is easy to open, and doing so does not void Apple's warranty. You merely need to remove one screw in the back of the case and pop off the top; instructions are given in the Mac II manual. Although you should be reasonably careful, a Mac II chassis is less hazardous than a small Mac's because all high-voltage supplies are in the video display, not in the Mac II. Furthermore, AC line-level voltages are confined within the power supply cage. If you want to make adjustments inside the video display, then the same cautions about opening a small Mac apply to the display unit.

ENVIRONMENT

Operating temperature: 50 to 104 degrees F (10 to 40 degrees C)

Storage temperature: −40 to 122 degrees F (−40 to 50 degrees C)

Altitude: sea level to 15,000 ft (4500 m) for small Macs, 10,000 ft (3000 m) for the Mac II

The temperatures at which you can run a Mac are limited largely by the floppy disks and floppy disk drives. Altitude limits result from an electrical property of air: The breakdown voltage, or voltage at which a spark can jump across a gap, decreases with decreasing air pressure and density. Thus, at high altitudes, the high-voltage power supply driving the CRT will arc. As a practical matter, this is likely to occur only at a few weather and research stations or in unpressurized aircraft. (Above the altitude limit, you can disable a small Mac's CRT and high-voltage power supply and use a more suitable monitor with a video-output jack, or you can use a flat-panel display.)

The Mac II is limited to lower altitudes than the small Macs because its larger video screens operate at a higher voltage. Similarly, the high-voltage coronas in the LaserWriters limit their operation to 8200 feet (2500 meters); at higher elevations you will have to use a conventional printer. These are operational limits; you can ship these units in unpressurized aircraft without problems.

LISA (MACINTOSH XL)

The Lisa was Apple's first visual-interface microcomputer. Like the small Macs, the Lisa used a 68000 CPU, but its hardware was organized differently and required a special systems program—MacWorks—to configure it to run Macintosh software. The Lisa used a 12-inch CRT with a 6¼-by-8¼-inch (16-by-21.5-centimeter) image area displaying 364 by 720 pixels. On the original Lisa, the pixels were oblong: 50 percent taller than their width (Macintosh pixels are square). Consequently, software that drew a perfect circle on the Macintosh drew an ellipse on the Lisa. Eventually the Lisa was renamed the Macintosh XL, and a minor adjustment in the video circuitry made the pixels more nearly square. The Macintosh XL was discontinued in 1985.

Glossary

accessory card: A hardware board (card), on which electronic circuitry has been printed, that adds features to a computer. Only the Macintosh SE and the Macintosh II are designed to accommodate internal accessory cards; the SE accepts one, the II accepts six. *See also* **expansion slot.**

acoustic modem: A modem with cups that fit snugly around the earpiece and mouthpiece of a standard telephone handset and contain a small speaker and microphone to convert a computer's digital signals into sound and back again. *See also* **direct-connect modem; modem.**

ADB: *See* **Apple DeskTop Bus.**

alert box: In Macintosh software, a warning on the screen to which you usually must respond before you can proceed.

algorithm: A step-by-step procedure for solving a problem or computing a result; specifically, the computer program code that does this.

analog-to-digital converter: An electronic circuit that changes continuous analog signals into discrete digital signals (bits); abbreviated A/D.

Apple DeskTop Bus (ADB): A simple local area network circuit designed to handle a keyboard, mouse, and other low-speed input devices; used in the Macintosh SE and Macintosh II (and the Apple IIGS).

Apple key: *See* **Command key (⌘).**

Apple menu: The menu under the Apple symbol at the left edge of the menu bar. *See also* **menu; menu bar.**

AppleTalk: A local area network developed by Apple Computer that can handle up to 32 nodes on a cable up to 1000 feet long.

application program: Software for accomplishing a specific task, such as writing (word processors), financial calculations (spreadsheets), and illustration (graphics programs). *See also* **utility program.**

arrow keys: A set of (usually) four keys found on many computer keyboards that move the cursor (insertion point) left, right, up, or down on the screen; also called cursor keys or direction keys.

ASCII (American Standard Code for Information Interchange): A standard scheme for coding characters in which each character is represented by seven bits; accepted by nearly all microcomputers and many large computers. *See also* **EBCDIC.**

assembler: A low-level programming language that is closely related to the intrinsic operation of a computer. Programs in assembler are fast and flexible but hard to read and write.

asynchronous communication: A means of transmitting data between computers in which a special signal indicates when each transmitted character starts and stops. *See also* **synchronous communication.**

audio output port: In the Macintosh, the electrical connection that transmits sound output; connected internally to a built-in speaker and externally to a ⅛-inch audio jack.

background program or process: A program or procedure that, once started, can continue without input from the user and without interfering with an active program running simultaneously in the foreground. *See also* **multitasking.**

BASIC (Beginner's All-purpose Symbolic Instruction Code): The most common programming language for microcomputers; relatively easy to learn but not especially flexible.

baud, baud rate: In digital communication, a measure of speed equal to one signal event per second. At low speeds (300 baud or less), one baud represents one bit per second; at higher speeds, one baud represents two or more bits. *See also* **bps.**

binary file: A computer file whose contents can be any arbitrary combination of bits. *See also* **text file.**

bit (contraction of **bi**nary digi**t**): The smallest unit of digital information, representing either an electronic on or off, coded as 1 or 0.

bit-mapped image: A screen or printed computer image consisting of dots, in which one dot in the image corresponds to one or more bits in the section of random-access memory set aside as a "map" of the screen.

boot: A computer's start-up process, the instructions for which are stored in ROM. The term comes from "pulling yourself up by your own bootstraps."

bps: Bits per second; a measure of data transfer speed. *See also* **baud.**

bridge: In a computer network, the hardware and software connecting two or more similar networks; all users on all connected networks have access to one another. *See also* **gateway.**

brownout: A period of low-voltage electrical power, usually caused by heavy demand or problems at the electrical utility.

buffer: A portion of memory that takes information from one device and feeds it to another; acts as a holding tank if the information is being sent from one device faster than it can be processed by the other.

bug: A software problem; named after a moth that caused the failure of an early (1945) digital computer at Harvard. Evocative but inaccurate entomology: Moths (and butterflies) belong to the order Lepidoptera, whereas true bugs belong to the order Hemiptera.

bulletin board, electronic: The computerized equivalent of a conventional notice board, maintained on a computer and available via modem and the phone line.

byte: A sequence of eight bits; usually represents one character.

cathode-ray tube (CRT): The screen used in nearly all televisions and most computers, in which light is produced by an electron beam (the cathode ray) striking a phosphor coating on the screen.

CCITT (Comité Consultatif Internationale de Télégraphie et Téléphonie): International Telegraph and Telephone Consultative Committee; an international standards-setting body, part of the International Telecommunications Union of the United Nations.

central processing unit (CPU): A computer's main information-processing circuit. In a microcomputer, the CPU is a single silicon chip called the microprocessor or CPU chip; on larger computers, the CPU may consist of many chips.

character width: A term used by Apple in its communication programs to refer to the number of bits, generally seven or eight, in a transmitted character; usually referred to as the number of data bits.

chip: A tiny wafer of silicon with a complex integrated electronic circuit photoengraved on its surface.

clicking: Pressing once and releasing the button(s) on a computer mouse.

Clipboard: A region of random-access memory that stores information copied or cut from a document you are working on with a Macintosh application. Clipboard information can be pasted into the same document, a different document created with the same application program, or a document created with a different program.

comma-delimited file: A data file in which commas separate data elements.

Command key (⌘): A special Macintosh shift key used in combination with other keys to issue commands, such as opening and closing files, cutting and pasting, and so on; also called the Apple key.

command-line interface: An old-fashioned way to get information into and out of a computer by means of cryptic commands displayed one line at a time on the screen.

communication protocol: A specific set of rules or procedures governing the way information is sent between computers.

compiler: A program that translates a high-level programming language (source code) into machine code that the computer understands; an entire program is generally compiled as a unit. *See also* **interpreter; source code.**

context switching: A type of processing in which several programs are loaded into memory simultaneously, but only one, the foreground program, is active at a time; background programs are frozen but can be brought into the foreground instantly. *See also* **background program; foreground program; multitasking.**

control character: A normally nonprinting ASCII character that controls operations or performs another function, such as coding for formatting in text; generated with the Command key on the Macintosh keyboards lacking a separate Control key.

copy-protected disk: A disk that cannot be copied, usually because it is in a nonstandard format.

CPU: *See* **central processing unit.**

CRT: *See* **cathode-ray tube.**

CSMA/CD (Carrier Sense Multiple Access with Collision Detection): A network protocol that permits any node to send information when the network cable is free; if two nodes start at the same time, each node pauses for a random interval before trying again.

cursor: The flashing marker indicating the current working location on a computer screen; called the insertion point on a Macintosh.

cursor keys: *See* **arrow keys.**

cylinder: On a disk, the aggregate of all the disk tracks that can be written or read for a specific head position. On a double-sided floppy disk, a cylinder consists of two tracks; on a hard disk, it consists of four or more tracks.

daisy wheel: The printing element in one type of letter-quality printer, which resembles a daisy with embossed letters at the tips of plastic "petals."

data: Information; the raw material processed with a computer program.

database: A structured collection of information, such as an address book, organized for storage, retrieval, and updating.

data bits: In computer communication, the bits that code for a transmitted character, usually a sequence of seven or eight bits; same as Apple's term **character width.**

data disk: A disk that contains only data, without programs or start-up information.

data fork: The portion of a Macintosh disk file generally containing the user's data. *See also* **header; resource fork.**

DCA (Document Content Architecture): IBM's format for transmitting text documents from one computer system to another; comes in two forms, **RFTDCA** (revisable) and **FFTDCA** (final form).

DCE (data communication equipment): A device on one side of an RS-232C serial communication link; the other side is **DTE** (data terminal equipment). DCE and DTE differ in the direction that signals travel through their connector pins; a DCE device can usually be converted to a DTE device with a change in connecting cables. The distinction is not made in the RS-422 ports on the Macintosh, which are wired identically.

desk accessory: A small convenience program in the Macintosh System file that can operate at the same time as an application program, for example, Alarm Clock, Control Panel, and so on.

Desktop: The image on the Macintosh screen showing disk icons and disk windows; produced by the Finder.

DeskTop file: A hidden data file created on all Macintosh disks by the Finder; contains housekeeping information that the Finder uses to keep track of file folders, icons, and so forth.

dialog box: In Macintosh software, a box on the screen asking for information that you must supply before the program can proceed.

digital circuit: An electronic circuit that works with information coded in binary digits.

digital-to-analog converter: An electronic circuit that changes discrete digital signals (bits) into continuous analog signals; abbreviated D/A.

digitizer pad: A computer accessory that registers the motion of a special stylus; used most often for drawing.

direct-connect modem: A modem that bypasses the telephone handset and plugs directly into a telephone line. *See also* **acoustic modem; modem.**

directory: A file stored on a disk that indexes the location of information on that disk.

disk buffer: A special portion of random-access memory that temporarily holds often-used information (for example, the directory) from a disk; speeds operations because the information does not have to be constantly exchanged with the disk.

disk controller: An electronic circuit that converts information on the microcomputer bus into a signal used by the disk drive heads to read and write on a disk.

disk drive: An electromechanical device that reads and writes on a magnetic disk.

disk drive port: In the Macintosh, a parallel port designed for connection to an external microfloppy disk drive. *See also* **parallel port.**

disk server: A disk drive on a network that is available to any user. Disk-server software divides a single disk into several logical volumes, which behave as if they were individual disks, and controls access to them.

DISOSS (DIstributed Office Support System): The standard information-exchange format developed by IBM for its office products and now used by other vendors; incorporates **DCA.** *See also* **SNA.**

dot-matrix printer: A printer that forms characters and graphics out of dots. Usually refers to an impact dot-matrix printer, in which the dots are pressed onto the paper by a pin pushing on an inked ribbon, but other printer types, including laser printers, also work in dot-matrix fashion.

dots per inch (dpi): A linear measure most often applied to screen and printer resolution.

dots per square inch (dpsi): An areal measure applied to screen and printer resolution.

double-clicking: Pointing at an object with a computer mouse and quickly pressing and releasing the mouse button twice.

dpi: *See* **dots per inch.**

dpsi: *See* **dots per square inch.**

dragging: Pointing at an object with a mouse, then pressing and holding down the mouse button while moving the mouse (and object) to another location on the screen.

driver: Software that manages a particular hardware item. For example, a keyboard driver interprets keystrokes into the form required by the computer.

DTE (data terminal equipment): *See* **DCE.**

duplex: In telecommunication, allowing two-way transmission. In full-duplex communication, such as conventional voice telephone calls, transmission occurs both ways simultaneously; in half-duplex, such as citizens-band radio, transmission can go only one way at a time.

Dvorak keyboard: A keyboard layout devised by August Dvorak that puts the most-often-used letters in English near the most powerful fingers.

EBCDIC (Extended Binary Coded Decimal Interchange Code): A standard scheme for coding characters in which each character is represented by a sequence of eight bits; used primarily by large computers. *See also* **ASCII.**

echo: In computer communication, characters returned to the sender by the receiving computer. The echo mirrors the original transmission.

electronic mail: A form of computer communication in which the sender transmits a message to a central computer, which stores the message until the recipient can retrieve it electronically.

emulation: The technique of getting one piece of hardware or software to mimic another.

encryption: Converting information into an unintelligible form for security, most often with a password that acts as a key for later decryption.

Enter key: A special key on the Macintosh keyboard, often used to complete a keyboard entry.

ergonomics: Considering the human element in engineering design. The standards are ill-defined; whether a device is "ergonomic" is usually decided by the advertising department.

Escape key: A special key on a computer keyboard, including the Apple Keyboard and Apple Extended Keyboard for the Macintosh SE and Macintosh II, which generates an "escape" code (ASCII 1B hex); needed by a Mac only when communicating with another type of computer.

expansion slot: A place inside the computer for connecting accessory circuit cards. The Macintosh SE has one, and the Macintosh II has six; earlier Macintosh models have none. *See also* **accessory card.**

exponentiation: Raising a number a to a power b; the number a is multiplied by itself b times. For example, 2^3 is $2 \times 2 \times 2$, or 2 raised to the power of 3.

facsimile (fax): A method for sending pictures over ordinary dial-up telephone lines.

FFTDCA (Final-Form-Text Document Content Architecture): A computer text-file format that contains the final page image. *See also* **DCA.**

field: A unit of information, such as an item in a database; a collection of related fields constitutes a record. In an address database, the street address and the zip code are fields; the entire address is a record. *See also* **record.**

file server: A node on a network that contains a disk drive, processor, and controlling software; available to any user. File-serving software controls access to individual files; multiuser software allows several users access to the same file simultaneously, although only one person at a time can make changes.

Finder: Macintosh system software that manages files and disk directories and creates the Desktop.

firmware: Programs embedded in a computer's circuitry; cannot be changed as easily as programs on disk (software) but are not as fixed as the other electronic circuits (hardware). Read-only memory programs are often called firmware.

FKEYs: Short utility programs that are started by pressing Command, Shift, and a number key from 1 through 0.

font: A complete set of characters in a given typeface and size, measured in points; includes uppercase and lowercase characters, numbers, punctuation marks, ligatures, common signs, and accents.

foreground program or process: The active program (and active window) when several programs are in memory at the same time. *See also* **multitasking.**

formant synthesis: A common method of synthesizing speech, based on a small number of resonant, or formant, frequencies.

formatting (a disk): Placing address markers on new disks so that the disk drive can locate information on them; same as **initializing.**

frame store or buffer: A memory buffer for video images.

gateway: In a computer network, the hardware and software connecting two or more dissimilar networks. *See also* **bridge.**

gigabyte (GB): 1024 megabytes; sometimes 1000 megabytes.

handshake: In computer communication, an electrical signal used by the receiving device to stop transmission from the sending device until the transmitted data can be processed.

hardware: The physical components of a computer—electronic parts, wires, cases, and so on.

Hayes-compatible modem: A modem that sets modes and features with a standard set of commands developed originally by Hayes Microcomputer Products.

header: The portion of a Macintosh disk file containing the file's directory and related information, including the name, type, and source of the file. *See also* **data fork; resource fork.**

hierarchical database: A database that organizes information in a tree-like structure.

high-level language: A programming language such as BASIC or Pascal that incorporates elements of English into its syntax.

I-beam: The standard Macintosh pointer for editing text.

icon: In the Macintosh, a small graphic symbol that represents a file or a function.

impact printer: A printer that forms characters or images by striking an inked ribbon against paper.

initializing (a disk): Placing address markers on new disks so that the drive can locate information; same as **formatting.**

ink-jet printer: A printer that forms characters or images by squirting tiny drops of ink onto paper.

integrated software: Software that can perform more than one task.

interface: The common boundary between two entities, such as user and computer or printer and computer.

internal modem: A modem built into a computer. Among Macintoshes, can only be installed in an SE or a II.

interpreter: A program that translates a high-level programming language into machine-readable code; the translation is done line by line. *See also* **compiler.**

ISDN (Integrated Services Digital Network): An internationally standardized telephone service based on digital signals.

KB: *See* **kilobyte.**

Kbps: *See* **kilobits per second.**

kerning: In phototypesetting, adjusting the spacing between letters; specifically, reducing the space between certain pairs of letters (kerning pairs) for the best appearance.

keyboard macro: A single set of keystrokes or a single command key that, when pressed, executes many keystrokes, mouse movements, or both; the user assigns macros with a keyboard redefinition program.

key disk: In a copy-protected program, the floppy disk you must use to start the program.

kilobits per second (Kbps): Thousands of bits per second; a measure of data transfer speed.

kilobyte (KB): 1024 bytes; the most common measure of computer file size or memory capacity. A typical double-spaced typewritten page is 1.5 KB.

laser printer: A printer that forms images by moving a laser beam across a photoconductive drum; after electrostatic development, the printer puts the image on paper with xerography.

light pen: A computer pointing device, shaped like a pen, in which a light receptor detects the electron beam moving across the face of a CRT.

linked files, live links: Interdependent documents organized so that when a change is made in one document, the dependent documents are also automatically changed.

local area network (LAN): A communication pathway linking computers and accessory devices over relatively short distances (usually within a few city blocks). *See also* **network; wide area network.**

logical volume: A software-defined partition in a single physical disk; to the user, each logical volume behaves as if it were an independent disk in an independent disk drive.

Macintosh, Charles: Scottish inventor (1766–1843) of waterproof rubberized fabric, used in mackintosh (with a *k*) rain jackets.

macro: *See* **keyboard macro.**

McIntosh, John: Discoverer and cultivator of the McIntosh apple in Ontario, 1796. His name was misspelled by Apple when the Macintosh project began.

McIntosh Laboratory: Manufacturer of high-fidelity equipment in Binghamton, NY; has licensed use of the Macintosh name to Apple Computer, Inc.

magnetic printer: A printer that uses an array of recording heads to create an image on a magnetic belt; toner containing magnetic particles develops the image, which is then electrostatically transferred to paper.

mainframe: A large traditional computer usually shared by many users. The central processing unit, disk drives, and tape drives are normally housed in an air-conditioned room; user-operated terminals may be scattered all over.

markup language: Coded formatting and printing commands that are embedded within ASCII text. For example, in the markup language TeX, x^2 is the marked-up form of x^2.

Mbps: *See* **megabits per second.**

megabits per second (Mbps): Millions of bits per second; a measure of data transfer speed.

megabyte (MB): 1024 kilobytes; sometimes 1000 kilobytes; a measure of file size and memory capacity.

menu: In software, a list of available programs, commands, or functions from which the user can choose.

menu bar: The horizontal strip at the top of the Macintosh screen that lists menu titles.

menu-driven interface: A boundary between user and computer in which the user issues commands by choosing items from a series of menus.

menu-initial interface: A boundary between user and computer in which the user issues commands by choosing menu items identified by a single letter.

menu-word interface: A boundary between user and computer in which the user issues commands by choosing menu items identified by single words arranged in a line or two on the screen.

microcomputer: A small computer designed in size and price to serve (primarily) a single person; also known as personal computer and home computer.

microfloppy (disk): A 3½-inch flexible disk within a semirigid plastic envelope; used in the Macintosh and many other computer brands.

microprocessor: A single silicon chip containing thousands of electronic components, capable of manipulating information when operated in conjunction with accessory devices.

mil: One thousandth of an inch.

millisecond (ms): One thousandth of a second.

minifloppy: A 5¼-inch flexible computer disk, very common in microcomputers.

modem (contraction of **mo**dulator/**dem**odulator)**:** An electronic circuit that converts digital signals into sound frequencies and back again for transmission over telephone lines.

moiré: The pattern created by the juxtaposition of two geometrically regular, repetitive structures (such as two sets of parallel lines or halftone screens).

mouse: Used with the Macintosh and other computers, a hand-operated pointing device that registers movement; a pointer on the computer screen moves in correspondence with the movements of the mouse on the user's desk.

ms: *See* **millisecond.**

MultiFinder: Macintosh system software that allows loading of multiple application programs into memory simultaneously and performs context switching among them.

multitasking: A computer processing technique in which several programs are loaded simultaneously into memory; the CPU pays attention to all the programs by switching rapidly among them in a procedure known as time slicing. *See also* **context switching; round-robin.**

nanosecond (ns): One billionth of a second (0.000000001 second).

native file: The original data file generated by an application program; usually cannot be directly read by another program.

network: An electronic hardware and software communication pathway linking multiple computers and accessories; any device can exchange information with any other device on the network.

node: Any device on a network that can send and/or receive information.

ns: *See* **nanosecond.**

null-modem: A communication cable between two computers; used in lieu of modems.

object code: In programming, the executable code or machine-language program produced by the compiler.

operating system: Essential software that controls the operation of a computer, directing information among components.

optical character recognition, optical character reader (OCR): The technology by which typed or printed documents are optically scanned and the text turned into the codes for characters that a computer can process; the device or software that does this. *See also* **scanner.**

optical disc: A disc with a reflective finish, on whose surface music or data is recorded in the form of tiny deformations and from which the music or data can be played back by a laser beam.

optical fiber: A long thin strand of glass that carries information as a modulated light beam; can handle far higher communication rates than wire connections.

Option key: A special Macintosh shift key used with other keys to produce nonstandard characters and symbols; similar to the Alternate key on other microcomputers.

overlay: A program fragment stored on disk until needed by the program core. Necessary with some large programs that cannot entirely fit into random-access memory; the overlays are brought in as needed, each replacing one no longer in use.

parallel port: An electrical connection to a computer through which eight or more bits are transmitted simultaneously in one direction; specifically, a Centronics parallel port.

parity bit: An extra bit tagged onto the end of a character or other unit of data, whose value is used to check for errors in the data. In computer communication, for example, even parity means that if the total number of 1s in a byte is even, then the parity bit is set to 1; if the number is odd, the bit is set to 0. If both sending and receiving devices come up with the same value after independently computing parity, the character is correct.

PBX (Private Branch Exchange): The private telephone switchboard within an organization.

phosphor: Any material that emits visible light when struck by an electron beam; used in CRTs.

pixel: A picture element, or dot, in an image. Sometimes elaborated into two types: a pel consisting only of on/off or black/white information and a pixel containing further attributes, such as shades of gray or color.

point: In typography, a unit of measurement approximately equal to 1/72 inch.

PostScript: A computer language developed by Adobe Systems, Inc., to describe an image—text as well as graphics—for printing; can also be adapted to driving a video display.

power-line conditioner: An electrical filter designed to remove potentially damaging surges and brief high-voltage transients; installs between the AC power supply and an electronic device.

printer buffer: A memory buffer between a computer and a printer that holds characters coming from the computer at high speed so that the printer can read them out more slowly.

printer driver: Software that converts printing information generated by an application program into the form required to drive a particular printer; also called a printing resource.

printer port: The serial port designed for connection to a printer. *See also* **serial port.**

printer server: A printer on a network; available to all network users.

printing resource: *See* **printer driver.**

program: The instructions that specify the operation of a central processor and other computer hardware; synonymous with software.

programming language: The words, symbols, numbers, and grammar used to give instructions to a computer.

proportional spacing: Printing in which wider letters (such as *M* or *W*) take up more space than narrow ones (such as *i* or *l*).

protocol: In telecommunications, a set of rules and procedures governing how information travels between computers and other electronic devices. Hardware protocols define parameters such as the timing and frequency of the electrical signal; software protocols specify the details of the signal's content.

pull-down menu: A computer menu that appears on the screen only when requested; until then, only the menu titles are visible. *See also* **menu.**

QuickDraw: The programs in the Macintosh ROM that generate images on the screen and for printers. Designed principally for video screen and dot-matrix printer resolution; does not offer the fine control of PostScript.

RAM: *See* **random-access memory.**

RAM disk: A portion of random-access memory set aside to behave like a disk drive.

random-access memory (RAM): An electronic storage area used by the computer to hold information it is currently working on.

raw data transfer rate: In a disk drive or on a network, the peak speed at which information is transferred to and from the computer; usually much higher than the average transfer rate.

read-only memory (ROM): An electronic storage area whose contents are set during manufacture and cannot be changed.

record: A set of items (fields) in a database. In an address database, a complete address is a record. *See also* **field.**

relational database: A database in which any field or record can be associated with any other field or record.

resource fork: The portion of a Macintosh disk file that contains the program code, font information, and other data not normally generated directly by the user. *See also* **data fork; header.**

RFTDCA (Revisable-Form-Text Document Content Architecture): A computer text-file format that contains both text and formatting information; used for converting documents from one word processor to another. *See also* **DCA.**

round-robin: A computer processing technique in which several programs are simultaneously loaded into memory, and the CPU pays attention to each one in turn.

RS-170: Recommended standard (RS) specification for a composite video signal compatible with broadcast standards in North America and Japan (straight video, not radio-frequency modulated); set by the Electronic Industries Association.

RS-422, 423, 232C: Recommended standards (RS) for serial computer interfaces; set by the Electronic Industries Association.

scanner: A device that scans a printed page and converts graphic images into a form that can be processed by a computer. *See also* **optical character recognition.**

Scrapbook: In the Macintosh, a means of saving and transferring information (text, pictures, or other data) among files created with different programs. The Scrapbook operates like the Clipboard except that it is saved on disk. *See also* **Clipboard.**

screen dump: A pixel-for-pixel screen image printed on paper or stored in a disk file.

screen font: A font designed for screen display and for printing by an ImageWriter or LaserWriter IISC.

scroll arrow, bar, box: On the Macintosh, the symbols along the borders of a window that show the window's position with respect to the document it is displaying and that allow the user to move through the document.

SCSI (Small Computer System Interface): A relatively high-speed bidirectional parallel port for communication between computer devices; most often used to connect a hard disk drive to a computer.

sector: On a disk, the smallest contiguous physical unit for recording information; several sectors make up a track.

sector reader: Software that can read and change disk sectors directly.

serial port: An electrical connection to a computer through which data are transmitted in series, one bit after another; specifically, an RS-232C or RS-422 port and, in the Macintosh, the printer and modem ports.

server: In a network, any device that can be shared by all users.

signal-to-noise-ratio (S/N): The ratio of the voltage of a received electric signal to the voltage of the interfering noise; usually given in decibels.

simplex: In telecommunication, allowing one-way transmission only, as in ordinary radio or television broadcasting.

small Macintosh: As used in this book, all the Macintosh models that are the same size as the original 1984 model; as of early 1988, included all models except the Macintosh II.

SNA (System Network Architecture): IBM's umbrella term for a wide range of protocols and standards for computer communications.

socket: In the AppleTalk network, a logical concept that routes data coming into a node to the correct application program, distinguishing an electronic mail program from a database, for example.

software: The instructions that direct the operation of a central processor and other computer hardware; programs, individually and collectively, are known as software.

source code or program: The original instructions (usually written in a high-level programming language) that an interpreter or compiler turns into machine code for execution on a computer.

spreadsheet: A program that manipulates values laid out in a rectangular grid; the user specifies interrelationships among the values, and the program can calculate results based on these relationships.

start-up disk, boot disk: A disk containing all the information necessary to start a computer.

stop bit: In asynchronous communication, one or more bits added to mark the end of each character.

structured data program: Any application program that stores information in a regular, defined way. A spreadsheet is a structured data program; a word processor or free-form graphics program is not.

Switcher: Macintosh system software that allows as many as four application programs to be loaded into RAM simultaneously; it has been superseded by MultiFinder.

SYLK (SYmbolic LinK): The file format developed by Microsoft for business applications such as spreadsheets and databases; used for transferring data between otherwise incompatible programs.

synchronous communication: A means of transmitting data between computers in blocks of many characters, with a timing (synchronizing) signal at the beginning of each block instead of start and stop signals for each character. *See also* **asynchronous communication.**

system disk: In the Macintosh, any disk containing the System file that can be used to start the computer.

tab-delimited file: A data file in which tabs separate data elements.

text file: A computer file that contains only sequences of bits that correspond to characters; also called an ASCII file. *See also* **binary file.**

thermal-transfer printer: A dot-matrix printer that uses small heated pins to melt small dots of pigment onto paper.

toner: The black plastic powder that functions as ink in photocopiers and laser printers.

touch pad: A computer pointing device that is operated by moving one's finger over a flat surface.

touch screen: A computer screen that allows the user to point at objects by touching the screen.

track: A path on magnetic media for information storage. On a disk, tracks are concentric circles on the surface, made up of sectors; one or more tracks make up a cylinder.

trackball: A computer pointing device with a large roller that the user turns.

uninterruptible power supply (UPS): A power system that protects against power failures by operating continuously from a rechargeable battery.

utility program: Software needed to support a computer's operation rather than a user application. *See also* **application program.**

V.22 bis: Modem protocol (data-link layer) for 2400-bits-per-second transmission over telephone lines; set by the CCITT. V.22 is a 1200-bps standard; the word *bis* is French for *a second time.*

video-controller circuit: An electronic circuit that takes digital information and creates the signals necessary for displaying that information on a CRT.

video card: A circuit board containing the video controller and related circuits that plugs into a computer to control a video display.

video digitizer: A hardware accessory that converts standard analog video signals into a digital form for computers.

video RAM: A portion of random-access memory set aside for screen information. In the Macintosh, video RAM stores a bit map of the screen display.

visual interface: A modern computer interface using icons and other visually symbolic information instead of pure text.

warm boot: The process of resetting a computer to its start-up state without shutting off the power.

wide area network: A communication pathway linking computers and accessory devices over cross-continental distances, for example, the Department of Defense's ARPANET (Advanced Research Projects Agency Network). *See also* **network; local area network.**

window: An area in which a document is displayed on the screen; can be controlled independently from any other window.

wordwrap: A text entry feature in word-processing software that automatically advances to the next line at the end of a line; words are preserved as units.

write-protect tab: A small notch or part built into a disk jacket that is covered or set to prevent accidentally erasing the disk contents.

xerography: The most widely used photocopying technology, based on a photoconductive drum and electrostatic image development.

zone: In AppleTalk, an individual network that is connected to other AppleTalk networks via bridges.

INDEX

References to illustrations and tables are in italics.

A

B

G

J

K

L

P

Q

R

S

U

V

W

X

Z

Born in Qingdao, China, Cary Lu grew up in California and received an A.B. in physics from the University of California, Berkeley, and a Ph.D. in biology from the California Institute of Technology. In addiiton to making short films for *Sesame Street,* Lu has been an associate producer for *NOVA* and science and technology advisor for *Children's Television Workshop.* He was the founding managing editor of *High Technology* magazine and is currently technology editor for *Inc.* magazine. Lu is a cofounder of General Information, Inc., which publishes databases in print and electronic form.

Ellen W. Chu holds an A.B. and a Ph.D. in biology from the University of California, Santa Cruz, and an M.S. in zoology from the University of Paris. She has worked as a lecturer in Environmental Studies at UCSC, as an assistant professor in The Writing Program at the Massachusetts Institute of Technology, and as the editor for *BioScience* magazine. Chu is currently a freelance editor and writer. Her most recent projects include a textbook on the sociology of science.

The manuscript for this book was written on a Macintosh II and a Macintosh SE with Microsoft Word, version 3.01.

Cover design adapted from original design by Tim Girvin
Interior text design by Darcie S. Furlan
Illustrations by Martin Easley, Rick Bourgoin, and Becky Geisler-Johnson
Principal production art by Becky Geisler-Johnson

Text composition by Microsoft Press in Garamond with display in Helvetica Italic, using the Magna composition system and the Linotronic 300 laser imagesetter.

Mac

"An especially good investment for the new Mac owner or someone contemplating purchasing [a Mac]."
Online Today

"The one Macintosh book you'd choose if you could have only one. Virtually anything you might want to know at any level is here."
Macazine

"If you have a Mac or are curious about them, this is the book to have."
BYTE

Revised by Cary Lu and Ellen W. Chu

Quoted by industry experts, recommended in classrooms, passed around in user groups, and relied upon by everyday Macintosh users, THE APPLE® MACINTOSH® BOOK is the bible of the Macintosh world.

If you already own a Mac® and are thinking of upgrading, this book will show you the exciting range of new possibilities that exist in the Macintosh world. If you are deciding whether to buy a Mac, this book provides a superb, all-around introduction to Macintosh technology and will help you do some comparison shopping.

This is the newest edition of the book that accompanied the introduction of the first Apple Macintosh in 1984. It's been thoroughly updated to include the latest information on the explosion of new and emerging developments in the Mac market. Microcomputer expert Cary Lu and science editor Ellen W. Chu provide a straightforward look at:

- how the Mac works: its video screen, keyboard and mouse, disks and disk drives, data pathways and memory, and printers
- what's currently and soon to be available in word processing, graphics, business software, communications, and desktop publishing for the Mac
- how to manage disks with available memory and what options exist for mass storage
- what programming languages exist for the Mac, including HyperTalk™, C, BASIC, Logo, and others
- how to work with local area networks and AppleTalk®; communication between Macs and PCs, or Macs and mainframes
- future directions for the Apple Macintosh

THE APPLE MACINTOSH BOOK. Your Apple Macintosh library is incomplete without it.

Microsoft PRESS®

U.S.A.	**$21.95**
U.K.	£19.95
Austral.	$32.95

(recommended)

ISBN 1-55615-110-1

9 781556 151101