

301/1590 -6

**MICROPROCESSOR LABORATORY SECOND CENTRAL
AMERICAN COURSE ON ADVANCED VLSI DESIGN TECHNIQUES**

29 November - 17 December 2004

Puebla - MEXICO

INTRODUCTION TO UNIX/LINUX

**Maria Liz CRESPO
ICTP
Trieste
Italy**

These lecture notes are intended for distribution to participants only

UNIX is a trademark of X/Open
MS-DOS and Microsoft Windows are trademarks of Microsoft Corporation
OS/2 and Operating System/2 are trademarks of IBM
X Window System is a trademark of X Consortium, Inc.
Motif is a trademark of the Open Software Foundation
LINUX is not a trademark, and has no connection to UNIX, Unix System Laboratories, or to X/Open.
Please bring all unacknowledged trademarks to the attention of the author.

The LINUX Users' Guide

Copyright © 1993, 1994, 1996 Larry Greenfield

All you need to know to start using LINUX, a free Unix clone. This manual covers the basic Unix commands, as well as the more specific LINUX ones. This manual is intended for the beginning Unix user, although it may be useful for more experienced users for reference purposes.

Copyright © Larry Greenfield
427 Harrison Avenue
Highland Park, NJ
08904
leg+@andrew.cmu.edu

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.






Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided also that the sections that reprint "The GNU General Public License", "The GNU Library General Public License", and other clearly marked sections held under separate copyright are reproduced under the conditions given within them, and provided that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language under the conditions for modified versions. "The GNU General Public License" and "The GNU Library General Public License" may be included in a translation approved by the Free Software Foundation instead of in the original English.

At your option, you may distribute verbatim and modified versions of this document under the terms of the GNU General Public License, excepting the clearly marked sections held under separate copyright.

Exceptions to these rules may be granted for various purposes: Write to Larry Greenfield at the above address or email leg+@andrew.cmu.edu, and ask. It is requested (but not required) that you notify the author whenever commercially or large-scale printing this document. Royalties and donations are accepted and will encourage further editions.

These are some of the typographical conventions used in this book.

Bold	Used to mark new concepts, WARNINGS , and keywords in a language.
<i>italics</i>	Used for <i>emphasis</i> in text.
<i>slanted</i>	Used to mark meta-variables in the text, especially in representations of the command line. For example, " <i>ls -l foo</i> " where <i>foo</i> would "stand for" a filename, such as <i>/bin/cp</i> .
Typewriter	Used to represent screen interaction. Also used for code examples, whether it is "C" code, a shell script, or something else, and to display general files, such as configuration files. When necessary for clarity's sake, these examples or figures will be enclosed in thin boxes.
	Represents a key to press. You will often see it in this form: "Press  to continue."
	A diamond in the margin, like a black diamond on a ski hill, marks "danger" or "caution." Read paragraphs marked this way carefully.
	This X in the margin indicates special instructions for users of the X Window System.
	This indicates a paragraph that contains special information that should be read carefully.

Acknowledgments

The author would like to thank the following people for their invaluable help either with LINUX itself, or in writing *The LINUX Users' Guide*:

Linus Torvalds for providing something to write this manual about.

Karl Fogel has given me much help with writing my LINUX documentation and wrote most of Chapter 8 and Chapter 9. I cannot give him enough credit.

Maurizio Codogno wrote much of Chapter 11.

David Channon wrote the appendix on *v1*. (Appendix A)

Yggdrasil Computing, Inc. for their generous (and voluntary) support of this manual.

Red Hat Software for their (more recent and still voluntary!) support.

The **fortune** program for supplying me with many of the wonderful quotes that start each chapter. They cheer me up, if no one else.

Chapter 2

What's Unix, anyway?

Ken Thompson has an automobile which he helped design. Unlike most automobiles, it has neither speedometer, nor gas gage, nor any of the numerous idiot lights which plague the modern driver. Rather, if the driver makes any mistake, a giant "?" lights up in the center of the dashboard. "The experienced driver," he says, "will usually know what's wrong."

2.1 Unix History

In 1965, Bell Telephone Laboratories (Bell Labs, a division of AT&T) was working with General Electric and Project MAC of MIT to write an operating system called Multics. To make a long story slightly shorter, Bell Labs decided the project wasn't going anywhere and broke out of the group. This left Bell Labs without a good operating system.

Ken Thompson and Dennis Ritchie decided to sketch out an operating system that would meet Bell Labs' needs. When Thompson needed a development environment (1970) to run on a PDP-7, he implemented their ideas. As a pun on Multics, Brian Kernighan, another Bell Labs researcher, gave the system the name Unix.

Later, Dennis Ritchie invented the "C" programming language. In 1973, Unix was rewritten in C instead of the original assembly language.¹ In 1977, Unix was moved to a new machine through a process called porting away from the PDP machines it had run on previously. This was aided by the fact Unix was written in C since much of the code could simply be recompiled and didn't have to be rewritten.

In the late 1970's, AT&T was forbidden from competing in the computing industry, so it licensed Unix to various colleges and universities very cheaply. It was slow to catch on outside of academic institutions but was eventually popular with businesses as well. The Unix of today is different from the Unix of 1970. It has two major variations: System V, from Unix System Laboratories

¹"Assembly language" is a very basic computer language that is tied to a particular type of computer. It is usually considered a challenge to program in.

(USL), a subsidiary of Novell², and the Berkeley Software Distribution (BSD). The USL version is now up to its fourth release, or SVR4³, while BSD's latest version is 4.4. However, there are many different versions of Unix besides these two. Most commercial versions of Unix derive from one of the two groupings. The versions of Unix that are actually used usually incorporate features from both variations.

Current commercial versions of Unix for Intel PCs cost between \$500 and \$2000.

2.2 LINUX History

The primary author of LINUX is Linus Torvalds. Since his original versions, it has been improved by countless numbers of people around the world. It is a clone, written entirely from scratch, of the Unix operating system. Neither USL, nor the University of California, Berkeley, were involved in writing LINUX. One of the more interesting facts about LINUX is that development occurs simultaneously around the world. People from Australia to Finland contributed to LINUX and will hopefully continue to do so.

LINUX began with a project to explore the 386 chip. One of Linus's earlier projects was a program that would switch between printing AAAA and BBBB. This later evolved to LINUX.

LINUX has been copyrighted under the terms of the GNU General Public License (GPL). This is a license written by the Free Software Foundation (FSF) that is designed to prevent people from restricting the distribution of software. In brief, it says that although you can charge as much as you'd like for a copy, you can't prevent the person you sold it to from giving it away for free. It also means that the source code⁴ must also be available. This is useful for programmers. Anybody can modify LINUX and even distributed his/her modifications, provided that they keep the code under the same copyright.

LINUX supports most of popular Unix software, including the X Window System. The X Window System was created at the Massachusetts Institute of Technology. It was written to allow Unix systems to create graphical windows and easily interact with each other. Today, the X Window System is used on every version of Unix available.

In addition to the two variations of Unix, System V and BSD, there is also a set of standardization documents published by the IEEE entitled POSIX. LINUX is first and foremost compliant with the POSIX-1 and POSIX-2 documents. Its look and feel is much like BSD in some places, and somewhat like System V in others. It is a blend (and to most people, a good one) of all three standards.

Many of the utilities included with LINUX distributions are from the Free Software Foundation and are part of GNU Project. The GNU Project is an effort to write a portable, advanced operating system that will look a lot like Unix. "Portable" means that it will run on a variety of machines, not just Intel PCs, Macintoshes, or whatever. The GNU Project's operating system is called the Hurd. The main difference between LINUX and GNU Hurd is not in the user interface but in the

²It was recently sold to Novell. Previously, USL was owned by AT&T.

³A cryptic way of saying "system five, release four".

⁴The source code of a program is what the programmer reads and writes. It is later translated into unreadable machine code that the computer interprets.

programmer's interface—the Hurd is a modern operating system while LINUX borrows more from the original Unix design.

The above history of LINUX is deficient in mentioning anybody *besides* Linux Torvalds. For instance, H. J. Lu has maintained `gcc` and the LINUX C Library (two items needed for all the programs on LINUX) since very early in LINUX's life. You can find a list of people who deserve to be recognized on every LINUX system in the file `/usr/src/linux/CREDITS`.

2.2.1 LINUX Now

The first number in LINUX's version number indicates truly huge revisions. These change very slowly and as of this writing (February, 1996) only version "1" is available. The second number indicates less major revisions. Even second numbers signify more stable, dependable versions of LINUX while odd numbers are developing versions that are more prone to bugs. The final version number is the minor release number—every time a new version is released that may just fix small problems or add minor features, that number is increased by one. As of February, 1996, the latest stable version is 1.2.11 and the latest development version is 1.3.61.

LINUX is a large system and unfortunately contains bugs which are found and then fixed. Although some people still experience bugs regularly, it is normally because of non-standard or faulty hardware; bugs that effect everyone are now few and far between.

Of course, those are just the kernel bugs. Bugs can be present in almost every facet of the system, and inexperienced users have trouble separating different programs from each other. For instance, a problem might arise that all the characters are some type of gibberish—is it a bug or a "feature"? Surprisingly, this is a feature—the gibberish is caused by certain control sequences that somehow appeared. Hopefully, this book will help you to tell the different situations apart.

2.2.2 A Few Questions and Answers

Before we embark on our long voyage, let's get the ultra-important out of the way.

Question: Just how do you pronounce LINUX?

Answer: According to Linus, it should be pronounced with a short *ih* sound, like `print`, `minimal`, etc. LINUX should rhyme with `Minix`, another Unix clone. It should *not* be pronounced like (American pronunciation of) the "Peanuts" character, `Linus`, but rather *LIH-nucks*. And the *u* is sharp as in `rule`, not soft as in `ducks`. LINUX should almost rhyme with "cynics".

Question: Why work on LINUX?

Answer: Why not? LINUX is generally cheaper (or at least no more expensive) than other operating systems and is frequently less problematic than many commercial systems. It might not be the best system for your particular applications, but for someone who is interested in using Unix applications available on LINUX, it is a high-performance system.

2.2.3 Commercial Software in LINUX

There is a lot of commercial software available for LINUX. Starting with `Motif`, a user interface for the X Window System that vaguely resembles Microsoft Windows, LINUX has been gaining more and more commercial software. These days you can buy anything from `Word Perfect` (a popular word processor) to `Maple`, a complex symbolic manipulation package, for LINUX.

For any readers interested in the legalities of LINUX, this is allowed by the LINUX license. While the GNU General Public License (reproduced in Appendix B) covers the LINUX kernel and would seemingly bar commercial software, the GNU Library General Public License (reproduced in Appendix C) covers most of the computer code applications depend on. This allows commercial software providers to sell their applications and withhold the source code.

Please note that those two documents are copyright notices, and not licenses to use. They do *not* regulate how you may use the software, merely under what circumstances you can copy it and any derivative works. To the Free Software Foundation, this is an important distinction: LINUX doesn't involve any "shrink-wrap" licenses but is merely protected by the same law that keeps you from photocopying a book.

Chapter 3

Getting Started

This login session: \$13.99, but for you \$11.88.

You may have previous experience with MS-DOS or other single user operating systems, such as OS/2 or the Macintosh. In these operating systems, you didn't have to identify yourself to the computer before using it; it was assumed that you were the only user of the system and could access everything. Well, Unix is a multi-user operating system—not only can more than one person use it at a time, different people are treated differently.

To tell people apart, Unix needs a user to identify him or herself¹ by a process called logging in. When you first turn on the computer a complex process takes place before the computer is ready for someone to use it. Since this guide is geared towards LINUX, I'll tell you what happens during the LINUX boot-up sequence.

If you're using LINUX on some type of computer besides an Intel PC, some things in this chapter won't apply to you. Mostly, they'll be in Section 3.1.

If you're just interested in using your computer, you can skip all the information in the chapter except for Section 3.3.

3.1 Power to the Computer


The first thing that happens when you turn an Intel PC on is that the BIOS executes. BIOS stands for Basic Input/Output System. It's a program permanently stored in the computer on read-only chips. It performs some minimal tests, and then looks for a floppy disk in the first disk drive. If it finds one, it looks for a "boot sector" on that disk, and starts executing code from it, if any. If there is a disk, but no boot sector, the BIOS will print a message like:

Non-system disk or disk error

¹From here on in this book, I shall be using the masculine pronouns to identify all people. This is the standard English convention, and people shouldn't take it as a statement that only men can use computers.

5

Figure 3.1 The path an Intel PC takes to get to a shell prompt. init may or may not start the X Window System. If it does, xdm runs. Otherwise, getty runs.



Removing the disk and pressing a key will cause the boot process to continue.

If there isn't a floppy disk in the drive, the BIOS looks for a master boot record (MBR) on the hard disk. It will start executing the code found there, which loads the operating system. On LINUX systems, LILO, the LINUX LOader, can occupy the MBR position, and will load LINUX. For now, we'll assume that happens and that LINUX starts to load. (Your particular distribution may handle booting from the hard disk differently. Check with the documentation included with the distribution. Another good reference is the LILO documentation, [1].)

3.2 LINUX Takes Over

After the BIOS passes control to LILO, LILO passes control to the LINUX kernel. A kernel is the central program of the operating system, in control of all other programs. The first thing that LINUX does once it starts executing is to change to protected mode. The 80386² CPU that controls your computer has two modes called "real mode" and "protected mode". DOS runs in real mode, as does the BIOS. However, for more advanced operating systems, it is necessary to run in protected mode. Therefore, when LINUX boots, it discards the BIOS.

Other CPUs will get to this stage differently. No other CPU needs to switch into protected mode and few have to have such a heavy framework around the loading procedure as LILO and the BIOS. Once the kernel starts up, LINUX works much the same.

LINUX then looks at the type of hardware it's running on. It wants to know what type of hard disks you have, whether or not you have a bus mouse, whether or not you're on a network, and other bits of trivia like that. LINUX can't remember things between boots, so it has to ask these questions each time it starts up. Luckily, it isn't asking you these questions—it is asking the hardware!

²When I refer to the 80386, I am also talking about the 80486, Pentium, and Pentium Pro computers unless I specifically say so. Also, I'll be abbreviating 80386 as 386.

During boot-up, the LINUX kernel will print variations on several messages. You can read about the messages in Section 3.4. This query process can cause problems with your system but if it was going to, it probably would have when you first installed LINUX. If you're having problems, consult your distribution's documentation.

The kernel merely manages other programs, so once it is satisfied everything is okay, it must start another program to do anything useful. The program the kernel starts is called `init`. (Notice the difference in font. Things in this font are usually the names of programs, files, directories, or other computer related items.) After the kernel starts `init`, it never starts another program. The kernel becomes a manager and a provider, not an active program.

So to see what the computer is doing after the kernel boots up, we'll have to examine `init`. `init` goes through a complicated startup sequence that isn't the same for all computers. LINUX has many different versions of `init`, and each does things its own way. It also matters whether your computer is on a network and what distribution you used to install LINUX. Some things that might happen once `init` is started:

- The file systems might be checked. What is a file system? A file system is the layout of files on the hard disk. It let's LINUX know which parts of the disk are already used, and which aren't. (It's like an index to a rather large filing system or a card catalog to a library.) Unfortunately, due to various factors such as power losses, what the file system information thinks is going on in the rest of the disk and the actual layout of the rest of the disk are occasionally in conflict. A special program, called `fsck`, can find these situations and hopefully correct them.
- Special routing programs for networks are run. These programs tell your computer how it's suppose to contact other computers.
- Temporary files left by some programs may be deleted.
- The system clock can be correctly updated. This is trickier then one might think, since Unix, by default, wants the time in UCT (Universal Coordinated Time, also known as Greenwich Mean Time) and your CMOS clock, a battery powered clock in your computer, is probably set on local time. This means that some program must read the time from your hardware clock and correct it to UCT.

After `init` is finished with its duties at boot-up, it goes on to its regularly scheduled activities. `init` can be called the parent of all processes on a Unix system. A process is simply a running program. Since one program can be running two or more times, there can be two or more processes for any particular program.

In Unix, a process, an instance of a program, is created by a system call—a service provided by the kernel—called `fork`. (It's called "fork" since one process splits off into two separate ones.) `init` forks a couple of processes, which in turn fork some of their own. On your LINUX system, what `init` runs are several instances of a program called `getty`. `getty` is the program that will allow a user to login and eventually calls a program called `login`.

3.3 The User Acts

3.3.1 Logging In

The first thing you have to do to use a Unix machine is to identify yourself. The login is Unix's way of knowing that users are authorized to use the system. It asks for an account name and password. An account name is normally similar to your regular name; you should have already received one from your system administrator, or created your own if you are the system administrator. (Information on doing this should be available in *Installation and Getting Started* or *The LINUX System Administrator's Guide*.)

You should see, after all the boot-up procedures are done, something like the following (the first line is merely a greeting message—it might be a disclaimer or anything else):

```
Welcome to the mousehouse. Please, have some cheese.
```

```
mousehouse login:
```



However, it's possible that what the system presents you with does *not* look like this. Instead of a boring text mode screen, it is graphical. However, it will still ask you to login, and will function mostly the same way. If this is the case on your system, you are going to be using The X Window System. This means that you will be presented with a windowing system. Chapter 5 will discuss some of the differences that you'll be facing. Logging in will be similar as will the basics to much of Unix. If you are using X, look for a giant X in the margin.

This is, of course, your invitation to login. Throughout this manual, we'll be using the fictional (or not so fictional, depending on your machine) user `larry`. Whenever you see `larry`, you should be substituting your own account name. Account names are usually based on real names; bigger, more serious Unix systems will have accounts using the user's last name, or some combination of first and last name, or even some numbers. Possible accounts for Larry Greenfield might be: `larry`, `greenfie`, `lgreenfi`, `lg19`.

`mousehouse` is, by the way, the "name" of the machine I'm working on. It is possible that when you installed LINUX, you were prompted for some very witty name. It isn't very important, but whenever it comes up, I'll be using `mousehouse` or, rarely, `lionsden` when I need to use a second system for clarity or contrast.

After entering `larry` and pressing `[return]`, I'm faced with the following:

```
mousehouse login: larry
Password:
```

What LINUX is asking for is your password. When you type in your password, you won't be able to see what you type. Type carefully: it is possible to delete, but you won't be able to see what you are editing. Don't type too slowly if people are watching—they'll be able to learn your password. If you mistype, you'll be presented with another chance to login.

If you've typed your login name and password correctly, a short message will appear, called the message of the day. This could say anything—the system administrator decides what it should

be. After that, a **prompt** appears. A prompt is just that, something prompting you for the next command to give the system. It should look something like this:

```
/home/larry#
```

X If you've already determined you're using X, you'll probably see a prompt like the one above in a "window" somewhere on the screen. (A "window" is a rectangular box.) To type into the prompt, move the mouse cursor (it probably looks like a big "x" or an arrow) using the mouse into the window.

3.3.2 Leaving the Computer

Do not just turn off the computer! You risk losing valuable data!

SLOW Unlike most versions of DOS, it's a bad thing to just hit the power switch when you're done using the computer. It is also bad to reboot the machine (with the reset button) without first taking proper precautions. LINUX, in order to improve performance, has a **disk cache**. This means it temporarily stores part of the computer's permanent storage in RAM.³ The idea of what LINUX thinks the disk should be and what the disk actually contains is synchronized every 30 seconds. In order to turn off or reboot the computer, you'll have to go through a procedure telling it to stop caching disk information.

If you're done with the computer, but are logged in (you've entered a username and password), first you must **logout**. To do so, enter the command **logout**. All commands are sent by pressing **return**. Until you hit return nothing will happen and you can delete what you've done and start over.

```
/home/larry# logout
```

```
Welcome to the mousehouse. Please, have some cheese.
```

```
mousehouse login:
```

Now another user can login.

3.3.3 Turning the Computer Off

If this is a single user system, you might want to turn the computer off when you're done with it.⁴ To do so, you'll have to log into a special account called **root**. The **root** account is the system administrator's account and can access any file on the system. If you're going to turn the computer

³The difference between "RAM" and a hard disk is like the difference between short term memory and long term memory. Shutting off the power is like giving the computer a knock on the head—it'll forget everything in short term memory. But things saved in long term memory, the hard disk, will be okay. The disk is thousands of times slower than RAM.

⁴To avoid possibly weakening some hardware components, only turn off the computer when you're done for the day. Turning the computer on and off once a day is probably the best compromise between energy and wear & tear on the system.

off, get the password from the system administrator. (In a single user system, that's *you!* Make sure you know the root password.) Login as root:

```
mousehouse login: root
Password:
Linux version 1.3.55 (root@mousehouse) #1 Sun Jan 7 14:56:26 EST 1996
/# shutdown now
Why? end of the day
```

```
URGENT: message from the sysadmin:
System going down NOW
```

```
... end of the day ...
```

Now you can turn off the power...

The command **shutdown now** prepares the system to be reset or turned off. Wait for a message saying it is safe to and then reset or turn off the system. (When the system asks you "Why?", it is merely asking for a reason to tell other users. Since no one is using the system when you shut it down, you can tell it anything you want or nothing at all.)

A quick message to the lazy: an alternative to the **logout/login** approach is to use the command **su**. As a normal user, from your prompt, type **su** and press **return**. It should prompt you for the root password, and then give you root privileges. Now you can **shutdown** the system with the **shutdown now** command.

3.4 Kernel Messages

When you first start your computer, a series of messages flash across the screen describing the hardware that is attached to your computer. These messages are printed by the LINUX kernel. In this section, I'll attempt to describe and explain those messages.

Naturally, these messages differ from machine to machine. I'll describe the messages I get for my machine. The following example contains all of the standard messages and some specific ones. (In general, the machine I'm taking this from is a minimally configured one: you won't see a lot of device specific configuration.) This was made with Linux version 1.3.55 - one of the most recent as of this writing.

1. The first thing LINUX does is decides what type of video card and screen you have, so it can pick a good font size. (The smaller the font, the more that can fit on the screen on any one time.) LINUX may ask you if you want a special font, or it might have had a choice compiled in.⁵

```
Console: 16 point font, 400 scans
```

```
Console: colour VGA+ 80x25, 1 virtual console (max 63)
```

⁵"Compiled" is the process by which a computer program that a human writes gets translated into something the computer understands. A feature that has been "compiled in" has been included in the program.

In this example, the machine owner decided he wanted the standard, large font at compile time. Also, note the misspelling of the word "color." Linus evidently learned the wrong version of English.

- The next thing the kernel will report is how fast your system is, as measured by "BogoMIPS". A "MIP" stands for a million instructions per second, and a "BogoMIP" is a "bogus MIP": how many times the computer can do absolutely nothing in one second. (Since this loop doesn't actually do anything, the number is not actually a measure of how fast the system is.) LINUX uses this number when it needs to wait for a hardware device.

```
Calibrating delay loop.. ok - 33.28 BogoMIPS
```

- The LINUX kernel also tells you a little about memory usage:

```
Memory: 23180k/24576k available (644k kernel code, 384k reserved, 468k data)
```

This said that the machine had 24 megabytes of memory. Some of this memory was reserved for the kernel. The rest of it can be used by programs. This is the temporary RAM that is used only for short term storage. Your computer also has a permanent memory called a hard disk. The hard disk's contents stay around even when power is turned off.

- Throughout the bootup procedure, LINUX tests different parts of the hardware and prints messages about these tests.

```
This processor honours the WP bit even when in supervisor mode. Good.
```

- Now LINUX moves onto the network configuration. The following should be described in *The LINUX Networking Guide*, and is beyond the scope of this document.

```
Swansea University Computer Society NET3.033 for Linux 1.3.50
IP Protocols: ICMP, UDP, TCP
```

- LINUX supports a FPU, a floating point unit. This is a special chip (or part of a chip, in the case of a 80486DX CPU) that performs arithmetic dealing with non-whole numbers. Some of these chips are bad, and when LINUX tries to identify these chips, the machine "crashes". The machine stops functioning. If this happens, you'll see:

```
Checking 386/387 coupling...
```

Otherwise, you'll see:

```
Checking 386/387 coupling... Ok, fpu using exception 16 error reporting.
```

if you're using a 486DX. If you are using a 386 with a 387, you'll see:

```
Checking 386/387 coupling... Ok, fpu using irq13 error reporting.
```

- It now runs another test on the "halt" instruction.

```
Checking 'hlt' instruction... Ok.
```

- After that initial configuration, LINUX prints a line identifying itself. It says what version it is, what version of the GNU C Compiler compiled it, and when it was compiled.

```
Linux version 1.3.55 (root@mousehouse) (gcc version 2.7.0) #1 Sun Jan 7 14:56:26 EST 1996
```

- The serial driver has started to ask questions about the hardware. A driver is a part of the kernel that controls a device, usually a peripheral. It is responsible for the details of how the CPU communicates with the device. This allows people who write user applications to concentrate on the application: they don't have to worry about exactly how the computer works.

```
Serial driver version 4.11 with no serial options enabled
tty00 at 0x03f8 (irq = 4) is a 16450
tty01 at 0x02f8 (irq = 3) is a 16450
tty02 at 0x03e8 (irq = 4) is a 16450
```

Here, it found 3 serial ports. A serial port is the equivalent of a DOS COM port, and is a device normally used to communicate with modems and mice.

What it is trying to say is that serial port 0 (COM1) has an address of 0x03f8. When it interrupts the kernel, usually to say that it has data, it uses IRQ 4. An IRQ is another means of a peripheral talking to the software. Each serial port also has a controller chip. The usual one for a port to have is a 16450; other values possible are 8250 and 16550.

- Next comes the parallel port driver. A parallel port is normally connected to a printer, and the names for the parallel ports (in LINUX) start with lp. lp stands for Line Printer, although in modern times it makes more sense for it to stand for Laser Printer. (However, LINUX will happily communicate with any sort of parallel printer: dot matrix, ink jet, or laser.)

```
lp0 at 0x03bc, (polling)
```

That message says it has found one parallel port, and is using the standard driver for it.

- LINUX next identifies your hard disk drives. In the example system I'm showing you, mousehouse, I've installed two IDE hard disk drives.

```
hda: WDC AC2340, 325MB w/127KB Cache, CHS=1010/12/55
hdb: WDC AC2850F, 814MB w/64KB Cache, LBA, CHS=827/32/63
```

- The kernel now moves onto looking at your floppy drives. In this example, the machine has two drives: drive "A" is a 5 1/4 inch drive, and drive "B" is a 3 1/2 inch drive. LINUX calls drive "A" fd0 and drive "B" fd1.

```
Floppy drive(s): fd0 is 1.44M, fd1 is 1.2M
floppy: FDC 0 is a National Semiconductor PC87306
```

- The next driver to start on my example system is the SLIP driver. It prints out a message about its configuration.

```
SLIP: version 0.8.3-NET3.019-HEWITY (dynamic channels, max=256) (6 bit encapsulation enabled)
CSLIP: code copyright 1989 Regents of the University of California
```

- The kernel also scans the hard disks it found. It will look for the different partitions on each of them. A partition is a logical separation on a drive that is used to keep operating systems from interfering with each other. In this example, the computer had two hard disks (hda, hdb) with four partitions and one partition, respectively.

```
Partition check:  
hda: hda1 hda2 hda3 hda4  
hdb: hdb1
```

15. Finally, LINUX mounts the root partition. The root partition is the disk partition where the LINUX operating system resides. When LINUX "mounts" this partition, it is making the partition available for use by the user.

```
VFS: Mounted root (ext2 filesystem) readonly.
```

Chapter 4

The Unix Shell

Making files is easy under the UNIX operating system. Therefore, users tend to create numerous files using large amounts of file space. It has been said that the only standard thing about all UNIX systems is the message-of-the-day telling users to clean up their files.
System V.2 administrator's guide

4.1 Unix Commands

When you first log into a Unix system, you are presented with something that looks like the following:

```
/home/larry#
```

That “something” is called a **prompt**. As its name would suggest, it is prompting you to enter a command. Every Unix command is a sequence of letters, numbers, and characters. There are no spaces, however. Some valid Unix commands are `mail`, `cat`, and `CHU.is.Number-5`. Some characters aren't allowed—we'll go into that later. Unix is also case-sensitive. This means that `cat` and `Cat` are different commands.¹

The prompt is displayed by a special program called the **shell**. Shells accept commands, and run those commands. They can also be programmed in their own language, and programs written in that language are called “shell scripts”.

There are two major types of shells in Unix: Bourne shells and C shells. Bourne shells are named after their inventor, Steven Bourne. Steven Bourne wrote the original Unix shell `sh`, and most shells since then end in the letters `sh` to indicate they are extensions on the original idea. There are many implementations of his shell, and all those specific shell programs are called Bourne shells. Another class of shells, C shells (originally implemented by Bill Joy), are also common. Traditionally, Bourne shells have been used for shell scripts and compatibility with the original `sh` while C shells have been

¹Case sensitivity is a very personal thing. Some operating systems, such as OS/2 or Windows NT are case preserving, but not case sensitive. In practice, Unix rarely uses the different cases. It is unusual to have a situation where `cat` and `Cat` are different commands.

used for interactive use. (C shells have had the advantages of having better interactive features but somewhat harder programming features.)

LINUX comes with a Bourne shell called `bash`, written by the Free Software Foundation. `bash` stands for Bourne Again Shell, one of the many bad puns in Unix. It is an “advanced” Bourne shell: it contains the standard programming features found in all Bourne shells with many interactive features commonly found in C shells. `bash` is the default shell to use running LINUX.

When you first login, the prompt is displayed by `bash`, and you are running your first Unix program, the `bash` shell. As long as you are logged in, the `bash` shell will constantly be running.

4.1.1 A Typical Unix Command

The first command to know is `cat`. To use it, type `cat`, and then `return`:

```
/home/larry# cat
```

If you now have a cursor on a line by itself, you've done the correct thing. There are several variances you could have typed—some would work, some wouldn't.

- If you misspelled `cat`, you would have seen

```
/home/larry# ct
ct: command not found
/home/larry#
```

Thus, the shell informs you that it couldn't find a program named “`ct`” and gives you another prompt to work with. Remember, Unix is case sensitive: `CAT` is a misspelling.

- You could have also placed whitespace before the command, like this:²

```
/home/larry#   cat
```

This produces the correct result and runs the `cat` program.

- You might also press `return` on a line by itself. Go right ahead—it does absolutely nothing.

I assume you are now in `cat`. Hopefully, you're wondering what it is doing. No, it is not a game. `cat` is a useful utility that won't seem useful at first. Type anything and hit `return`. What you should have seen is:

```
/home/larry# cat
Help! I'm stuck in a Linux program!
Help! I'm stuck in a Linux program!
```

²The ‘ ’ indicates that the user typed a space.

(The *slanted* text indicates what I typed to `cat`.) What `cat` seems to do is echo the text right back at yourself. This is useful at times, but isn't right now. So let's get out of this program and move onto commands that have more obvious benefits.

To end many Unix commands, type `Ctrl-d`³. `Ctrl-d` is the end-of-file character, or EOF for short. Alternatively, it stands for end-of-text, depending on what book you read. I'll refer to it as an end-of-file. It is a control character that tells Unix programs that you (or another program) is done entering data. When `cat` sees you aren't typing anything else, it terminates.

For a similar idea, try the program `sort`. As its name indicates, it is a sorting program. If you type a couple of lines, then press `Ctrl-d`, it will output those lines in a sorted order. These types of programs are called filters, because they take in text, filter it, and output the text slightly differently. Both `cat` and `sort` are unusual filters. `cat` is unusual because it reads in text and performs no changes on it. `sort` is unusual because it reads in lines and doesn't output anything until after it's seen the EOF character. Many filters run on a line-by-line basis: they will read in a line, perform some computations, and output a different line.

4.2 Helping Yourself

The `man` command displays reference pages for the command⁴ you specify. For example:

```
/home/larry# man cat

cat(1)                                cat(1)

NAME
  cat - Concatenates or displays files

SYNOPSIS
  cat [-benstuvAET] [--number] [--number-nonblank] [--squeeze-blank]
  [--show-nonprinting] [--show-ends] [--show-tabs] [--show-all]
  [--help] [--version] [file...]

DESCRIPTION
  This manual page documents the GNU version of cat ...
```

There's about one full page of information about `cat`. Try running `man` now. Don't expect to understand the manpage given. Manpages usually assume quite a bit of Unix knowledge—knowledge that you might not have yet. When you've read the page, there's probably a little black block at the bottom of your screen similar to “--more--” or “Line 1”. This is the *more-prompt*, and you'll learn to love it.

³Hold down the key labeled “Ctrl” and press “d”, then let go of both.

⁴`man` will also display information on a system call, a subroutine, a file format, and more. In the original version of Unix it showed the exact same information the printed documentation would. For now, you're probably only interested in getting help on commands.

Instead of just letting the text scroll away, `man` stops at the end of each page, waiting for you to decide what to do now. If you just want to go on, press `Space` and you'll advance a page. If you want to exit (quit) the manual page you are reading, just press `q`. You'll be back at the shell prompt, and it'll be waiting for you to enter a new command.

There's also a keyword function in `man`. For example, say you're interested in any commands that deal with Postscript, the printer control language from Adobe. Type `man -k ps` or `man -k Postscript`, you'll get a listing of all commands, system calls, and other documented parts of Unix that have the word “ps” (or “Postscript”) in their name or short description. This can be very useful when you're looking for a tool to do something, but you don't know its name—or if it even exists!

4.3 Storing Information

Filters are very useful once you are an experienced user, but they have one small problem. How do you store the information? Surely you aren't expected to type everything in each time you are going to use the program! Of course not. Unix provides files and directories.

A directory is like a folder: it contains pieces of paper, or files. A large folder can even hold other folders—directories can be inside directories. In Unix, the collection of directories and files is called the file system. Initially, the file system consists of one directory, called the “root” directory. Inside this directory, there are more directories, and inside those directories are files and yet more directories.⁵


Each file and each directory has a name. It has both a short name, which can be the same as another file or directory somewhere else on the system, and a long name which is unique. A short name for a file could be `joë`, while its “full name” would be `/home/larry/joë`. The full name is usually called the path. The path can be decoded into a sequence of directories. For example, here is how `/home/larry/joë` is read:

```
/home/larry/joë
The initial slash indicates the root directory.
This signifies the directory called home. It is inside the root directory.
  This is the directory larry, which is inside home.
    joë is inside larry. A path could refer to either a directory or a filename,
    so joë could be either. All the items before the short name must be directories.
```

An easy way of visualizing this is a tree diagram. To see a diagram of a typical LINUX system, look at Figure 4.1. Please note that this diagram isn't complete—a full LINUX system has over 8000 files!—and shows only some of the standard directories. Thus, there may be some directories in that diagram that aren't on your system, and your system almost certainly has directories not listed there.

⁵There may or may not be a limit to how “deep” the file system can go. (I've never reached it—one can easily have directories 10 levels deep.)

Figure 4.1 A typical (abridged) Unix directory tree.



4.3.1 Looking at Directories with ls

Now that you know that files and directories exist, there must be some way of manipulating them. Indeed there is. The command `ls` is one of the more important ones. It lists files. If you try `ls` as a command, you'll see:

```
/home/larry# ls
/home/larry#
```

That's right, you'll see nothing. Unix is intentionally terse: it gives you nothing, not even "no files" if there aren't any files. Thus, the lack of output was `ls`'s way of saying it didn't find any files.

But I just said there could be 8000 or more files lying around: where are they? You've run into the concept of a "current" directory. You can see in your prompt that your current directory is `/home/larry`, where you don't have any files. If you want a list of files of a more active directory, try the root directory:

```
/home/larry# ls /
bin      etc      install  mnt      root     user     var
dev      home    lib      proc     tmp      usr      vmlinux
/home/larry#
```

In the above command, "`ls /`", the directory ("`/`") is a parameter. The first word of the command is the command name, and anything after it is a parameter. Parameters generally modify what the program is acting on—for `ls`, the parameters say what directory you want a list for. Some commands have special parameters called options or switches. To see this try:

```
/home/larry# ls -F /
bin/      etc/      install/  mnt/      root/     user/     var@
dev/      home/    lib/      proc/     tmp/      usr/      vmlinux
/home/larry#
```

The `-F` is an option. An option is a special kind of parameter that starts with a dash and modifies how the program runs, but not what the program runs on. For `ls`, `-F` is an option that lets you see which ones are directories, which ones are special files, which are programs, and which are normal files. Anything with a slash is a directory. We'll talk more about `ls`'s features later. It's a surprisingly complex program!

Now, there are two lessons to be learned here. First, you should learn what `ls` does. Try a few other directories that are shown in Figure 4.1, and see what they contain. Naturally, some will be empty, and some will have many, many files in them. I suggest you try `ls` both with and without the `-F` option. For example, `ls /usr/local` looks like:

```
/home/larry# ls /usr/local
archives  bin      emacs  etc      ka9q  lib      tcl
/home/larry#
```

The second lesson is more general. Many Unix commands are like `ls`. They have options, which are generally one character after a dash, and they have parameters. Unlike `ls`, some commands require certain parameters and/or options. To show what commands generally look like, we'll use the following form:

```
ls [-aRF] [directory]
```

I'll generally use command templates like that before I introduce any command from now on. The first word is the command (in this case `ls`). Following the command are all the parameters. Optional parameters are contained in brackets ("`[`" and "`]`"). Meta-variables are *slanted*—they're words that take the place of actual parameters. (For example, above you see `directory`, which should be replaced by the name of a real directory.)

Options are a special case. They're enclosed by brackets, but you can take any one of them without using all of them. For instance, with just the three options given for `ls` you have eight different ways of running the command: with or without each of the options. (Contrast `ls -R` with `ls -F`.)

4.3.2 The Current Directory and cd

pwd

Using directories would be cumbersome if you had to type the full path each time you wanted to access a directory. Instead, Unix shells have a feature called the “current” or “present” or “working” directory. Your setup most likely displays your directory in your prompt: `/home/larry`. If it doesn't, try the command `pwd`, for present working directory. (Sometimes the prompt will display the machine name. This is only really useful in a networked environment with lots of different machines.)

```
mousehouse>pwd
/home/larry
mousehouse>
```

cd [directory]

As you can see, `pwd` tells you your current directory⁶—a very simple command. Most commands act, by default, on the current directory. For instance, `ls` without any parameters displays the contents of the current directory. We can change our current directory using `cd`. For instance, try:

```
/home/larry# cd /home
/home# ls -F
larry/  sam/  shutdown/  steve/  user1/
/home#
```

If you omit the optional parameter *directory*, you're returned to your home, or original, directory. Otherwise, `cd` will change you to the specified directory. For instance:

```
/home# cd
/home/larry# cd /
/# cd home
/home# cd /usr
/usr# cd local/bin
/usr/local/bin#
```

As you can see, `cd` allows you to give either absolute or relative pathnames. An absolute path starts with `/` and specifies all the directories before the one you wanted. A relative path is in relation to your current directory. In the above example, when I was in `/usr`, I made a relative move to `local/bin`—`local` is a directory under `usr`, and `bin` is a directory under `local`! (`cd home` was also a relative directory change.)

⁶You'll see all the terms in this book: present working directory, current directory, or working directory. I prefer “current directory”, although at times the other forms will be used for stylistic purposes.

There are two directories used *only* for relative pathnames: “.” and “..”. The directory “.” refers to the current directory and “..” is the parent directory. These are “shortcut” directories. They exist in *every* directory, but don't really fit the “folder in a folder” concept. Even the root directory has a parent directory – it's its own parent!

The file `./chapter-1` would be the file called `chapter-1` in the current directory. Occasionally, you need to put the “./” for some commands to work, although this is rare. In most cases, `./chapter-1` and `chapter-1` will be identical.

The directory “..” is most useful in “backing up”:

```
/usr/local/bin# cd ..
/usr/local# ls -F
archives/  bin/      emacs0    etc/      ks9q/     lib/      tc10
/usr/local# ls -F ../src
cweb/     linux/    xmr1s/
/usr/local#
```

In this example, I changed to the parent directory using `cd ..`, and I listed the directory `/usr/src` from `/usr/local` using `../src`. Note that if I was in `/home/larry`, typing `ls -F ../src` wouldn't do me any good!

The directory `~/` is an alias for your home directory:

```
/usr/local# ls -F ~/
/usr/local#
```

You can see at a glance that there isn't anything in your home directory! `~/` will become more useful as we learn more about how to manipulate files.

4.3.3 Creating and Removing Directories

`mkdir directory1 [directory2 ... directoryN]`

Creating your own directories is extremely simple under Unix, and can be a useful organizational tool. To create a new directory, use the command `mkdir`. Of course, `mkdir` stands for make directory.

Let's do a small example to see how this works:

```
/home/larry# ls -F
/home/larry# mkdir report-1993
/home/larry# ls -F
report-1993/
/home/larry# cd report-1993
/home/larry/report-1993#
```

`mkdir` can take more than one parameter, interpreting each parameter as another directory to create. You can specify either the full pathname or a relative pathname; `report-1993` in the above example is a relative pathname.

```
/home/larry/report-1993# mkdir /home/larry/report-1993/chap1 ~/report-1993/chap2
/home/larry/report-1993# ls -F
chap1/ chap2/
/home/larry/report-1993#
```

`rmdir` *directory1* [*directory2* ... *directoryN*]

The opposite of `mkdir` is `rmdir` (remove directory). `rmdir` works exactly like `mkdir`.

An example of `rmdir` is:

```
/home/larry/report-1993# rmdir chap1 chap3
rmdir: chap3: No such file or directory
/home/larry/report-1993# ls -F
chap2/
/home/larry/report-1993# cd ..
/home/larry# rmdir report-1993
rmdir: report-1993: Directory not empty
/home/larry#
```

As you can see, `rmdir` will refuse to remove a non-existent directory, as well as a directory that has anything in it. (Remember, `report-1993` has a subdirectory, `chap2`, in it!) There is one more interesting thing to think about `rmdir`: what happens if you try to remove your current directory? Let's find out:

```
/home/larry# cd report-1993
/home/larry/report-1993# ls -F
chap2/
/home/larry/report-1993# rmdir chap2
/home/larry/report-1993# rmdir .
rmdir: .: Operation not permitted
/home/larry/report-1993#
```

Another situation you might want to consider is what happens if you try to remove the parent of your current directory. This turns out not to be a problem since the parent of your current directory isn't empty, so it can't be removed!

4.4 Moving Information

All of these fancy directories are very nice, but they really don't help unless you have some place to store your data. The Unix Gods saw this problem, and they fixed it by giving the users files.

We will learn more about creating and editing files in the next few chapters.

The primary commands for manipulating files under Unix are `cp`, `mv`, and `rm`. They stand for copy, move, and remove, respectively.

4.4.1 cp Like a Monk

`cp` [-i] *source destination*
`cp` [-i] *file1 file2 ... fileN destination-directory*⁷

`cp` is a very useful utility under Unix, and extremely powerful. It enables one person to copy more information in a second than a fourteenth century monk could do in a year.

Be careful with `cp` if you don't have a lot of disk space. No one wants to see a "Disk full" message when working on important files. `cp` can also overwrite existing files without warning—I'll talk more about that danger later.

We'll first talk about the first line in the command template. The first parameter to `cp` is the file to copy—the second is where to copy it. You can copy to either a different filename, or a different directory. Let's try some examples:

```
/home/larry# ls -F /etc/passwd
/etc/passwd
/home/larry# cp /etc/passwd .
/home/larry# ls -F
passwd
/home/larry# cp passwd frog
/home/larry# ls -F
frog passwd
/home/larry#
```

The first `cp` command I ran took the file `/etc/passwd`, which contains the names of all the users on the Unix system and their (encrypted) passwords, and copied it to my home directory. `cp` doesn't delete the source file, so I didn't do anything that could harm the system. So two copies of `/etc/passwd` exist on my system now, both named `passwd`, but one is in the directory `/etc` and one is in `/home/larry`.

Then I created a *third* copy of `/etc/passwd` when I typed `cp passwd frog`—the three copies are now: `/etc/passwd`, `/home/larry/passwd` and `/home/larry/frog`. The contents of these three files are the same, even if the names aren't.

`cp` can copy files between directories if the first parameter is a file and the second parameter is a directory. In this case, the short name of the file stays the same.

⁷`cp` has two lines in its template because the meaning of the second parameter can be different depending on the number of parameters.



It can copy a file and change its name if both parameters are file names. Here is one danger of `cp`. If I typed `cp /etc/passwd /etc/group`, `cp` would normally create a new file with the contents identical to `passwd` and name it `group`. However, if `/etc/group` already existed, `cp` would destroy the old file without giving you a chance to save it! (It won't even print out a message reminding you that you're destroying a file by copying over it.)

Let's look at another example of `cp`:

```
/home/larry# ls -F
frog passwd
/home/larry# mkdir passwd_version
/home/larry# cp frog passwd passwd_version
/home/larry# ls -F
frog      passwd      passwd_version/
/home/larry# ls -F passwd_version
frog passwd
/home/larry#
```

How did I just use `cp`? Evidently, `cp` can take *more* than two parameters. (This is the second line in the command template.) What the above command did is copied all the files listed (`frog` and `passwd`) and placed them in the `passwd_version` directory. In fact, `cp` can take any number of parameters, and interprets the first $n - 1$ parameters to be files to copy, and the n^{th} parameter as what directory to copy them too.

You cannot rename files when you copy more than one at a time—they always keep their short name. This leads to an interesting question. What if I type `cp frog passwd toad`, where `frog` and `passwd` exist and `toad` isn't a directory? Try it and see.



4.4.2 Pruning Back with `rm`

```
rm [-i] file1 file2 ... fileN
```

Now that we've learned how to create millions of files with `cp` (and believe me, you'll find new ways to create more files soon), it may be useful to learn how to delete them. Actually, it's very simple: the command you're looking for is `rm`, and it works just like you'd expect: any file that's a parameter to `rm` gets deleted.

For example:

```
/home/larry# ls -F
frog      passwd      passwd_version/
/home/larry# rm frog toad passwd
rm: toad: No such file or directory
/home/larry# ls -F
passwd_version/
/home/larry#
```

As you can see, `rm` is extremely unfriendly. Not only does it not ask you for confirmation, but it will also delete things even if the whole command line wasn't correct. This could actually be dangerous. Consider the difference between these two commands:

```
/home/larry# ls -F
toad frog/
/home/larry# ls -F frog
toad
/home/larry# rm frog/toad
/home/larry#
```

and this

```
/home/larry# rm frog toad
rm: frog is a directory
/home/larry# ls -F
frog/
/home/larry#
```

As you can see, the difference of *one* character made a world of difference in the outcome of the command. It is vital that you check your command lines before hitting `return`!



4.4.3 A Forklift Can Be Very Handy

```
mv [-i] old-name new-name
mv [-i] file1 file2 ... fileN new-directory
```

Finally, the other file command you should be aware of is `mv`. `mv` looks a lot like `cp`, except that it deletes the original file after copying it. It's a lot like using `cp` and `rm` together. Let's take a look at what we can do:

```
/home/larry# cp /etc/passwd .
/home/larry# ls -F
passwd
/home/larry# mv passwd frog
/home/larry# ls -F
frog
/home/larry# mkdir report
/home/larry# mv frog report
/home/larry# ls -F
report/
/home/larry# ls -F report
frog
/home/larry#
```

As you can see, `mv` will rename a file if the second parameter is a file. If the second parameter is a directory, `mv` will move the file to the new directory, keeping its shortname the same.



You should be very careful with `mv`—it doesn't check to see if the file already exists, and will remove any old file in its way. For instance, if I had a file named `frog` already in my directory `report`, the command `mv frog report` would delete the file `~/report/frog` and replace it with `~/frog`.

In fact, there is one way to make `rm`, `cp` and `mv` ask you before deleting files. All three of these commands accept the `-i` option, which makes them query the user before removing any file. If you use an alias, you can make the shell do `rm -i` automatically when you type `rm`. You'll learn more about this later in Section 9.1.3 on page 90.

Chapter 5

The X Window System

The nice thing about standards is that there are so many of them to choose from.

Andrew S. Tanenbaum



This chapter only applies to those using the X Window System. If you encounter a screen with multiple windows, colors, or a cursor that is only movable with your mouse, you are using X. (If your screen consists of white characters on a black background, you are not currently using X. If you want to start it up, take a look at Section 5.1.)

5.1 Starting and Stopping the X Window System

5.1.1 Starting X

Even if X doesn't start automatically when you login, it is possible to start it from the regular text-mode shell prompt. There are two possible commands that will start X, either `startx` or `xinit`. Try `startx` first. If the shell complains that no such command is found, try using `xinit` and see if X starts. If neither command works, you may not have X installed on your system—consult local documentation for your distribution.

If the command runs but you are eventually returned to the black screen with the shell prompt, X is installed but not configured. Consult the documentation that came with your distribution on how to setup X.

5.1.2 Exiting X

Depending on how X is configured, there are two possible ways you might have to exit X. The first is if your window manager controls whether or not X is running. If it does, you'll have to exit X using a menu (see Section 5.4.8 on page 43). To display a menu, click a button on the background.

The important menu entry should be "Exit Window Manager" or "Exit X" or some entry containing the word "Exit". Try to find that entry (there could be more than one menu—try different mouse buttons!) and choose it.

The other method would be for a special `xterm` to control X. If this is the case, there is probably a window labeled "login" or "system xterm". To exit from X, move the mouse cursor into that window and type "exit".

If X was automatically started when you logged in, one of these methods should log you out. Simply login again to return. If you started X manually, these methods should return you to the text mode prompt. (If you wish to logout, type `logout` at this prompt.)

5.2 What is The X Window System?

The X Window System is a distributed, graphical method of working developed primarily at the Massachusetts Institute of Technology. It has since been passed to a consortium of vendors (aptly named "The X Consortium") and is being maintained by them.

The X Window System (hereafter abbreviated as "X"¹) has new versions every few years, called releases. As of this writing, the latest revision is X11R6, or release six. The eleven in X11 is officially the version number but there hasn't been a new version in many years, and one is not currently planned.

There are two terms when dealing with X that you should be familiar. The **client** is a X program. For instance, `xterm` is the client that displays your shell when you log on. The **server** is a program that provides services to the client program. For instance, the server draws the window for `xterm` and communicates with the user.

Since the client and the server are two separate programs, it is possible to run the client and the server *on two physically separate machines*. In addition to supplying a standard method of doing graphics, you can run a program on a remote machine (across the country, if you like!) and have it display on the workstation right in front of you.

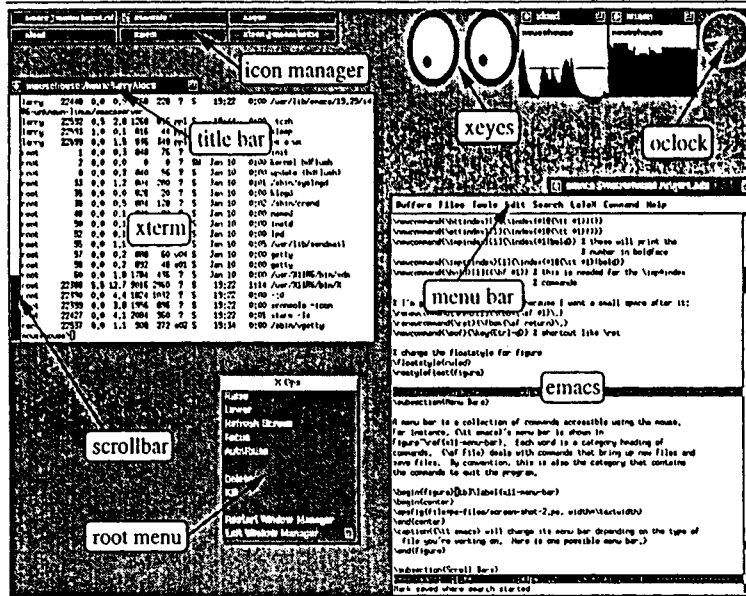
A third term you should be familiar with is the **window manager**. The window manager is a special client that tells the server where to position various windows and provides a way for the user to move these windows around. The server, by itself, does nothing for the user. It is merely there to provide a buffer between the user and the client.

5.3 What's This on my Screen?

When you first start X, several programs are started. First, the server is started. Then, several clients are usually started. Unfortunately, this is not standardized across various distributions. It is likely that among these clients are a window manager, either `fvwm` or `twm`, a prompt, `xterm`, and a clock, `xclock`.

¹There are several acceptable ways to refer to The X Window System. A common though incorrect way of referring to X is "X Windows".

Figure 5.1 An annotated example of a standard X screen. In this example, the user is running `twm`. The standard clock has been replaced by a transparent clock called `oclock`.



5.3.1 XClock

```
xclock [-digital] [-analog] [-update seconds] [-hands color]
```

I'll explain the simplest one first: `xclock` functions exactly as you'd expect it would. It ticks off the seconds, minutes and hours in a small window.

No amounts of clicking or typing in `xclock`'s window will affect it—that's *all* it does. Or is it? In fact, there are various different options you can give to the program to have it act in different ways. For instance, `xclock -digital` will create a digital clock. `xclock -update 1` will create a second hand that moves every second, while `-update 5` will create a second hand that moves every 5 seconds.

For more information on `xclock`'s options, consult its manpage—`man xclock`. If you're going to try running a few of your own `xclocks`, you should probably read Section 6.4 (Multitasking) to learn

how to run them in addition to your current programs. (If you run an `xclock` in the foreground—the usual way of running a program—and want to get out of it, type `ctrl-c`.)

5.3.2 XTerm

The window with a prompt in it (something that probably looks like `/home/larry#`) is being controlled by a program called `xterm`. `xterm` is a deceptively complicated program. At first glance, it doesn't seem to do much, but it actually has to do a lot of work. `xterm` emulates a terminal so that regular text-mode Unix applications work correctly. It also maintains a buffer of information so that you can refer back to old commands. (To see how to use this, look at Section 5.6.3.)

For much of this book, we're going to be learning about the Unix command-line, and you'll find that inside your `xterm` window. In order to type into `xterm`, you *usually* have to move your mouse cursor (possibly shaped like an "X" or an arrow) into the `xterm` window. However, this behavior is dependent on the window manager.

One way of starting more programs under X is through an `xterm`. Since X programs are standard Unix programs, they can be run from normal command prompts such as `xterms`. Since running a long term program from a `xterm` would tie up the `xterm` as long as the program was running, people normally start X programs in the background. For more information about this, see Section 6.4.

5.4 Window Managers

On LINUX, there are two different window managers that are commonly used. One of them, called `twm` is short for "Tab Window Manager". It is larger than the other window manager usually used, `fvwm`. (`fvwm` stands for "F(?) Virtual Window Manager"—the author neglected to tie down exactly what the `f` stood for.) Both `twm` and `fvwm` are highly configurable, which means I can't tell you exactly what keys do what in your particular setup.

To learn about `twm`'s configuration, look at Section 9.2.1. `fvwm`'s configuration is covered in Section 9.2.2.

5.4.1 When New Windows are Created

There are three possible things a window manager will do when a new window is created. It is possible to configure a window manager so that an outline of the new window is shown, and you are allowed to position it on your screen. That is called *manual placement*. If you are presented with the outline of a window, simply use the mouse to place it where you wish it to appear and click the left mouse button.

It is also possible that the window manager will place the new window somewhere on the screen by itself. This is known as *random placement*.

Finally, sometimes an application will ask for a specific spot on the screen, or the window manager will be configured to display certain applications on the same place of the screen all the time. (For

instance, I specify that I want `xclock` to always appear in the upper right hand corner of the screen.)

5.4.2 Focus

The window manager controls some important things. The first thing you'll be interested in is focus. The focus of the server is which window will get what you type into the keyboard. Usually in X the focus is determined by the position of the mouse cursor. If the mouse cursor is in one `xterm`'s window², that `xterm` will get your keypresses. This is different from many other windowing systems, such as Microsoft Windows, OS/2, or the Macintosh, where you must click the mouse in a window before that window gets focus. Usually under X, if your mouse cursor wanders from a window, focus will be lost and you'll no longer be able to type there.

Note, however, that it is possible to configure both `twm` and `fvwm` so that you must click on or in a window to gain focus, and click somewhere else to lose it, identical to the behavior of Microsoft Windows. Either discover how your window manager is configured by trial and error, or consult local documentation.

5.4.3 Moving Windows

Another very configurable thing in X is how to move windows around. In my personal configuration of `twm`, there are three different ways of moving windows around. The most obvious method is to move the mouse cursor onto the title bar and drag the window around the screen. Unfortunately, this may be done with any of the left, right, or middle buttons³. (To drag, move the cursor above the title bar, and hold down on the button while moving the mouse.) Most likely, your configuration is set to move windows using the *left* mouse buttons.

Another way of moving windows may be holding down a key while dragging the mouse. For instance, in my configuration, if I hold down the `Alt` key, move the cursor above a window, I can drag the window around using the left mouse button.

Again, you may be able to understand how the window manager is configured by trial and error, or by seeing local documentation. Alternatively, if you want to try to interpret the window manager's configuration file, see Section 9.2.1 for `twm` or Section 9.2.2 for `fvwm`.

5.4.4 Depth

Since windows are allowed to overlap in X, there is a concept of depth. Even though the windows and the screen are both two dimensional, one window can be in front of another, partially or completely obscuring the rear window.

There are several operations that deal with depth:

²You can have more than one copy of `xterm` running at the same time!

³Many PCs have only two button mice. If this is the case for you, you should be able to emulate a middle button by using the left and right buttons simultaneously.

- Raising the window, or bringing a window to the front. This is usually accomplished by clicking on a window's title bar with one of the buttons. Depending on how the window manager is configured, it could be any one of the buttons. (It is also possible that more than one button will do the job.)
- Lowering the window, or pushing the window to the back. This can generally be accomplished by a different click in the title bar. It is also possible to configure some window managers so that one click will bring the window forward if there is anything over it, while that same click will lower it when it is in the front.
- Cycling through windows is another operation many window managers allow. This brings each window to the front in an orderly cycle.

5.4.5 Iconization

There are several other operations that can obscure windows or hide them completely. First is the idea of "iconization". Depending on the window manager, this can be done in many different ways. In `twm`, many people configure an icon manager. This is a special window that contains a list of all the other windows on the screen. If you click on a name (depending on the setup, it could be with any of the buttons!) the window disappears—it is iconified. The window is still active, but you can't see it. Another click in the icon manager restores the window to the screen.

This is quite useful. For instance, you could have remote `xterms` to many different computers that you occasionally use. However, since you rarely use all of them at a given time, you can keep most of the `xterm` windows iconified while you work with a small subset. The only problem with this is it becomes easy to "lose" windows. This causes you to create new windows that duplicate the functionality of iconified windows.

Other window managers might create actual icons across the bottom of the screen, or might just leave icons on the root window.

5.4.6 Resizing

There are several different methods to resize windows under X. Again, it is dependent on your window manager and exactly how your window manager is configured. The method many Microsoft Windows users are familiar with is to click on and drag the border of a window. If your window manager creates large borders that change how the mouse cursor looks when it is moved over them, that is probably the method used to resize windows.

Another method used is to create a "resizing" button on the titlebar. In Figure 5.3, a small button is visible on the right of each titlebar. To resize windows, the mouse is moved onto the resize button and the left mouse button is held down. You can then move the mouse outside the borders of the window to resize it. The button is released when the desired size has been reached.

5.4.7 Maximization

Most window managers support maximization. In *twm*, for instance, you can maximize the height, the width, or both dimensions of a window. This is called "zooming" in *twm*'s language although I prefer the term maximization. Different applications respond differently to changes in their window size. (For instance, *xterm* won't make the font bigger but will give you a larger workspace.)

Unfortunately, it is extremely non-standard on how to maximize windows.

5.4.8 Menus

Another purpose for window managers is for them to provide menus for the user to quickly accomplish tasks that are done over and over. For instance, I might make a menu choice that automatically launches *Emacs* or an additional *xterm* for me. That way I don't need to type in an *xterm* an especially good thing if there aren't any running *xterms* that I need to type in to start a new program!

In general, different menus can be accessed by clicking on the root window, which is an immovable window behind all the other ones. By default, it is colored gray, but could look like anything.⁴ To try to see a menu, click and hold down a button on the desktop. A menu should pop up. To make a selection, move (without releasing the mouse button) the cursor over one of the items and then release the mouse button.

5.5 X Attributes

There are many programs that take advantage of X. Some programs, like *emacs*, can be run either as a text-mode program or as a program that creates its own X window. However, most X programs can only be run under X.

5.5.1 Geometry

There are a few things common to all programs running under X. In X, the concept of geometry is where and how large a window is. A window's geometry has four components:

- The horizontal size, usually measured in pixels. (A pixel is the smallest unit that can be colored. Many X setups on Intel PCs have 1024 pixels horizontally and 768 pixels vertically.) Some applications, like *xterm* and *emacs*, measure their size in terms of number of characters they can fit in the window. (For instance, eighty characters across.)
- The vertical size, also usually measured in pixels. It's possible for it to be measured in characters.

⁴One fun program to try is called *xfishtank*. It places a small aquarium in the background for you.

- The horizontal distance from one of the sides of the screen. For instance, +35 would mean make the left edge of the window thirty-five pixels from the left edge of the screen. On the other hand, -50 would mean make the right edge of the window fifty pixels from the right edge of the screen. It's generally impossible to start the window off the screen, although a window can be moved off the screen. (The main exception is when the window is very large.)
- The vertical distance from either the top or the bottom. A positive vertical distance is measured from the top of the screen; a negative vertical distance is measured from the bottom of the screen.

All four components get put together into a geometry string that looks like: 503x73-78+0. (That translates into a window 503 pixels long, 73 pixels high, put near the top right hand corner of the screen.) Another way of stating it is *hsize*x*vsize*±*hplace*±*vplace*.

5.5.2 Display

Every X application has a display that it is associated with. The display is the name of the screen that the X server controls. A display consists of three components:

- The machine name that the server is running on. At stand-alone LINUX installations the server is always running on the same system as the clients. In such cases, the machine name can be omitted.
- The number of the server running on that machine. Since any one machine could have multiple X servers running on it (unlikely for most LINUX machines, but possible) each must have a unique number.
- The screen number. X supports a particular server controlling more than one screen at a time. You can imagine that someone wants a lot of screen space, so they have two monitors sitting next to each other. Since they don't want two X servers running on one machine for performance reasons, they let one X server control both screens.

These three things are put together like so: *machine:server-number.screen-number*.

For instance, on *mousehouse*, all my applications have the display set to :0.0, which means the first screen of the first server on the local display. However, if I am using a remote computer, the display might be set to *mousehouse:0.0*.

By default, the display is taken from the environment variable (see Section 9.1.4) named *DISPLAY*, and can be overridden with a command-line option (see Figure 5.2). To see how *DISPLAY* is set, try the command *echo \$DISPLAY*.

5.6 Common Features

While X is a graphical user interface, it is a very uneven graphical user interface. It's impossible to say how any component of the system is going to work, because every component can easily be

Figure 5.2 Standard options for X programs.

Name	Followed by	Example
<code>-geometry</code>	geometry of the window	<code>xterm -geometry 80x24+0+90</code>
<code>-display</code>	display you want the program to appear	<code>xterm -display lionsden:0.0</code>
<code>-fg</code>	the primary foreground color	<code>xterm -fg yellow</code>
<code>-bg</code>	the primary background color	<code>xterm -bg blue</code>

reconfigured, changed, and even replaced. This means it's hard to say exactly how to use various parts of the interface. We've already encountered one cause of this: the different window managers and how configurable each window manager is.

Another cause of this uneven interface is the fact that X applications are built using things called "widget sets". Included with the standard X distribution are "Athena widgets" developed at MIT. These are commonly used in free applications. They have the disadvantage that they are not particularly good-looking and are somewhat harder to use than other widgets.

The other popular widget set is called "Motif". Motif is a commercial widget set similar to the user interface used in Microsoft Windows. Many commercial applications use Motif widgets, as well as some free applications. The popular World Wide Web Browser *netscape* uses Motif.

Let's try to go through some of the more usually things you'll encounter.

5.6.1 Buttons

Buttons are generally the easiest thing to use. A button is invoked by positioning the mouse cursor over it and clicking (pressing and immediately releasing the mouse button) the left button. Athena and Motif buttons are functionally the same although they have cosmetic differences.

5.6.2 Menu Bars

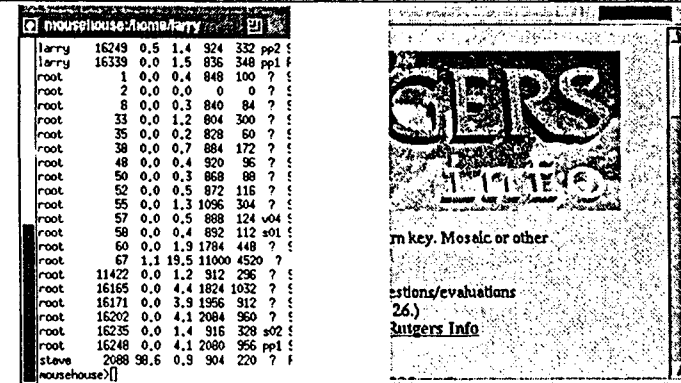
A menu bar is a collection of commands accessible using the mouse. For instance, *emacs*'s menu bar is shown in Figure 5.3. Each word is a category heading of commands. File deals with commands that bring up new files and save files. By convention, this is also the category that contains the command to exit the program.

To access a command, move the mouse cursor over a particular category (such as File) and press and hold down the left mouse button. This will display a variety of commands. To select one of the commands, move the mouse cursor over that command and release the left mouse button. Some menu bars let you click on a category—if this is the case, clicking on the category will display the menu until you click on either a command, another menu, or outside the menu bar (indicating that you are not interested in running a particular command).

Figure 5.3 *emacs* will change its menu bar depending on the type of file you're working on. Here is one possible menu bar.

```
Buffers Files Tools Edit Search Help
```

Figure 5.4 An Athena-type scroll bar is visible on the left of this *xterm* window. Next to it, a Motif-type scroll bar is visible on the *netscape* window.



5.6.3 Scroll Bars

A scroll bar is a method to allow people to display only part of a document, while the rest is off the screen. For instance, the *xterm* window is currently displaying the bottom third of the text available in Figure 5.4. It's easy to see what part of the available text is current being displayed: the darkened part of the scroll bar is relative to both the position and the amount of displayed text. If the text displayed is all there is, the entire scroll bar is dark. If the middle half of the text is displayed, the middle half of the scroll bar is darkened.

A vertical scroll bar may be to the left or right of the text and a horizontal one may be above or below, depending the application.

Athena scroll bars

Athena scroll bars operate differently from scroll bars in other windowing systems. Each of the three buttons of the mouse operate differently. To scroll upwards (that is, display material above what is currently visible) you can click the rightmost mouse button anywhere in the scroll bar. To scroll downwards, click the left mouse button anywhere in the scroll bar.

You can also jump to a particular location in the displayed material by clicking the middle mouse button anywhere in the scroll bar. This causes the window to display material starting at that point in the document.

Motif scroll bars

A Motif scroll bar acts much more like a Microsoft Windows or Macintosh scroll bar. An example of one is on the right in Figure 5.4. Notice that in addition to the bar, it has arrows above and below it. These are used for fine-tuning: clicking either the left or middle buttons on them will scroll a small amount such as one line; the right button does nothing.

The behavior of clicking inside the scroll bar is widely different for Motif scroll bars than Athena scroll bars. The right button has no effect. Clicking the left button above the current position scrolls upward. Similarly, clicking below the current position scrolls downward. Clicking and holding the left button on the current position allows one to move the bar at will. Releasing the left button positions the window.

Clicking the middle button anywhere on the bar will immediately jump to that location, similar to the behavior of the Athena middle button. However, instead of starting to display the data at the position clicked, that position is taken to be the *midpoint* of the data to be displayed.

Chapter 6

Working with Unix

A UNIX saleslady, Lenore,
 Enjoys work, but she likes the beach more.
 She found a good way
 To combine work and play:
 She sells C shells by the seashore.

Unix is a powerful system for those who know how to harness its power. In this chapter, I'll try to describe various ways to use Unix's shell, bash, more efficiently.

6.1 Wildcards

In the previous chapter, you learned about the file maintenance commands `cp`, `mv`, and `rm`. Occasionally, you want to deal with more than one file at once – in fact, you might want to deal with many files at once. For instance, you might want to copy all the files beginning with `data` into a directory called `~/backup`. You could do this by either running many `cp` commands, or you could list every file on one command line. Both of these methods would take a long time, however, and you have a large chance of making an error.

A better way of doing that task is to type:

```
/home/larry/report# ls -F
1993-1      1994-1      data1      data5
1993-2      data-new    data2
/home/larry/report# mkdir ~/backup
/home/larry/report# cp data* ~/backup
/home/larry/report# ls -F ~/backup
data-new    data1      data2      data5
/home/larry/report#
```

As you can see, the asterisk told `cp` to take all of the files beginning with `data` and copy them to `~/backup`. Can you guess what `cp d*w ~/backup` would have done?

6.1.1 What Really Happens?

Good question. Actually, there are a couple of special characters intercepted by the shell, `bash`. The character `*`, an asterisk, says "replace this word with all the files that will fit this specification". So, the command `cp data* ~/backup`, like the one above, gets changed to `cp data-new data1 data2 data5 ~/backup` before it gets run.

To illustrate this, let me introduce a new command, `echo`. `echo` is an extremely simple command; it echoes back, or prints out, any parameters. Thus:

```
/home/larry# echo Hello!
Hello!
/home/larry# echo How are you?
How are you?
/home/larry# cd report
/home/larry/report# ls -F
1993-1      1994-1      data1      data5
1993-2      data-new    data2
/home/larry/report# echo 199*
1993-1 1993-2 1994-1
/home/larry/report# echo *4*
1994-1
/home/larry/report# echo *2*
1993-2 data2
/home/larry/report#
```

As you can see, the shell expands the wildcard and passes all of the files to the program you tell it to run. This raises an interesting question: what happens if there are *no* files that meet the wildcard specification? Try `echo /rc/fr*og` and `bash` passes the wildcard specification verbatim to the program.

Other shells, like `tcsh`, will, instead of just passing the wildcard verbatim, will reply `No match`. Here's the same command run under `tcsh`:

```
mousehouse>echo /rc/fr*og
echo: No match.
mousehouse>
```

The last question you might want to know is what if I wanted to have `data*` echoed back at me, instead of the list of file names? Well, under both `bash` and `tcsh`, just include the string in quotes:

```
/home/larry/report# echo "data*"      mousehouse>echo "data*"
data*                                  OR      data*
/home/larry/report#                    mousehouse>
```

6.1.2 The Question Mark

In addition to the asterisk, the shell also interprets a question mark as a special character. A question mark will match one, and only one character. For instance, `ls /etc/??` will display all two letter files in the `/etc` directory.

6.2 Time Saving with bash

6.2.1 Command-Line Editing

Occasionally, you've typed a long command to `bash` and, before you hit return, notice that there was a spelling mistake early in the line. You could just delete all the way back and retype everything you need to, but that takes much too much effort! Instead, you can use the arrow keys to move back there, delete the bad character or two, and type the correct information.

There are many special keys to help you edit your command line, most of them similar to the commands used in GNU Emacs. For instance, `[C-t]` flips two adjacent characters.¹ You'll be able to find most of the commands in the chapter on Emacs, Chapter 8.

6.2.2 Command and File Completion

Another feature of `bash` is automatic completion of your command lines. For instance, let's look at the following example of a typical `cp` command:

```
/home/larry# ls -F
this-is-a-long-file
/home/larry# cp this-is-a-long-file shorter
/home/larry# ls -F
shorter      this-is-a-long-file
/home/larry#
```

It's a big pain to have to type every letter of `this-is-a-long-file` whenever you try to access it. So, create `this-is-a-long-file` by copying `/etc/passwd` to it². Now, we're going to do the above `cp` command very quickly and with a smaller chance of mistyping.

Instead of typing the whole filename, type `cp th` and press and release the `[Tab]`. Like magic, the rest of the filename shows up on the command line, and you can type in `shorter`. Unfortunately, `bash` cannot read your thoughts, and you'll have to type all of `shorter`.

When you type `[Tab]`, `bash` looks at what you've typed and looks for a file that starts like that. For instance, if I type `/usr/bin/ema` and then hit `[Tab]`, `bash` will find `/usr/bin/emacs` since that's the only file that begins `/usr/bin/ema` on my system. However, if I type `/usr/bin/ld` and hit `[Tab]`, `bash` beeps at me. That's because three files, `/usr/bin/ld`, `/usr/bin/ldd`, and `/usr/bin/ld86` all start with `/usr/bin/ld` on my system.

If you try a completion and `bash` beeps, you can immediately hit `[Tab]` again to get a list of all the files your start matches so far. That way, if you aren't sure of the exact spelling of your file, you can start it and scan a much smaller list of files.

¹`[C-t]` means hold down the key labeled "Ctrl", then press the "t" key. Then release the "Ctrl" key.

²`cp /etc/passwd this-is-a-long-file`

6.3 The Standard Input and The Standard Output

Let's try to tackle a simple problem: getting a listing of the `/usr/bin` directory. If all we do is `ls /usr/bin`, some of the files scroll off the top of the screen. How can we see all of the files?

6.3.1 Unix Concepts

The Unix operating system makes it very easy for programs to use the terminal. When a program writes something to your screen, it is using something called **standard output**. Standard output, abbreviated as `stdout`, is how the program writes things to a user. The name for what you tell a program is **standard input** (`stdin`). It's possible for a program to communicate with the user without using standard input or output, but most of the commands I cover in this book use `stdin` and `stdout`.

For example, the `ls` command prints the list of the directories to standard output, which is normally "connected" to your terminal. An interactive command, such as your shell, `bash`, reads your commands from standard input.

It is also possible for a program to write to **standard error**, since it is very easy to make standard output point somewhere besides your terminal. Standard error (`stderr`) is almost always connected to a terminal so an actual human will read the message.

In this section, we're going to examine three ways of fiddling with the standard input and output: input redirection, output redirection, and pipes.

6.3.2 Output Redirection

A very important feature of Unix is the ability to **redirect output**. This allows you, instead of viewing the results of a command, to save it in a file or send it directly to a printer. For instance, to redirect the output of the command `ls /usr/bin`, we place a `>` sign at the end of the line, and say what file we want the output to be put in:

```
/home/larry# ls
/home/larry# ls -F /usr/bin > listing
/home/larry# ls
listing
/home/larry#
```

As you can see, instead of writing the names of all the files, the command created a totally new file in your home directory. Let's try to take a look at this file using the command `cat`. If you think back, you'll remember `cat` was a fairly useless command that copied what you typed (the standard input) to the terminal (the standard output). `cat` can also print a file to the standard output if you list the file as a parameter to `cat`:

```
/home/larry# cat listing
...
/home/larry#
```

The exact output of the command `ls /usr/bin` appeared in the contents of `listing`. All well and good, although it didn't solve the original problem.³

However, `cat` does do some interesting things when its output is redirected. What does the command `cat listing > newfile` do? Normally, the `> newfile` says "take all the output of the command and put it in `newfile`." The output of the command `cat listing` is the file `listing`. So we've invented a new (and not so efficient) method of copying files.

How about the command `cat > fox`? `cat` by itself reads in each line typed at the terminal (standard input) and prints it right back out (standard output) until it reads `Ctrl-d`. In this case, standard output has been redirected into the file `fox`. Now `cat` is serving as a rudimentary editor:

```
/home/larry# cat > fox
The quick brown fox jumps over the lazy dog.
press Ctrl-d
```

We've now created the file `fox` that contains the sentence "The quick brown fox jumps over the lazy dog." One last use of the versatile `cat` command is to concatenate files together. `cat` will print out every file it was given as a parameter, one after another. So the command `cat listing fox` will print out the directory listing of `/usr/bin`, and then it will print out our silly sentence. Thus, the command `cat listing fox > listandfox` will create a new file containing the contents of both `listing` and `fox`.

6.3.3 Input Redirection

Like redirecting standard output, it is also possible to redirect standard input. Instead of a program reading from your keyboard, it will read from a file. Since input redirection is related to output redirection, it seems natural to make the special character for input redirection be `<`. It too, is used after the command you wish to run.

This is generally useful if you have a data file and a command that expects input from standard input. Most commands also let you specify a file to operate on, so `<` isn't used as much in day-to-day operations as other techniques.

6.3.4 The Pipe

Many Unix commands produce a large amount of information. For instance, it is not uncommon for a command like `ls /usr/bin` to produce more output than you can see on your screen. In order for you to be able to see all of the information that a command like `ls /usr/bin`, it's necessary to use another Unix command, called `more`.⁴ `more` will pause once every screenful of information. For instance, `more < /etc/rc` will display the file `/etc/rc` just like `cat /etc/rc` would, except that

³For impatient readers, the command you might want to try is `more`. However, there's still a bit more to talk about before we get there.

⁴`more` is named because that's the prompt it originally displayed: `--more--`. In many versions of LINUX the `more` command is identical to a more advanced command that does all that `more` can do and more. Proving that computer programmers make bad comedians, they named this new program `less`.

`more` will let you read it. `more` also allows the command `more /etc/rc`, and that's the normal way of invoking it.

However, that doesn't help the problem that `ls /usr/bin` displays more information than you can see. `more < ls /usr/bin` won't work—input redirection only works with files, not commands! You *could* do this:

```
/home/larry# ls /usr/bin > temp-ls
/home/larry# more temp-ls
...
/home/larry# rm temp-ls
```

However, Unix supplies a much cleaner way of doing that. You can just use the command `ls /usr/bin | more`. The character `|` indicates a pipe. Like a water pipe, a Unix pipe controls flow. Instead of water, we're controlling the flow of information!

A useful tool with pipes are programs called filters. A filter is a program that reads the standard input, changes it in some way, and outputs to standard output. `more` is a filter—it reads the data that it gets from standard input and displays it to standard output one screen at a time, letting you read the file. `more` isn't a great filter because its output isn't suitable for sending to another program.

Other filters include the programs `cat`, `sort`, `head`, and `tail`. For instance, if you wanted to read only the first ten lines of the output from `ls`, you could use `ls /usr/bin | head`.

6.4 Multitasking

6.4.1 Using Job Control

Job control refers to the ability to put processes (another word for programs, essentially) in the background and bring them to the foreground again. That is to say, you want to be able to make something run while you go and do other things, but have it be there again when you want to tell it something or stop it. In Unix, the main tool for job control is the shell—it will keep track of jobs for you, if you learn how to speak its language.

The two most important words in that language are `fg`, for foreground, and `bg`, for background. To find out how they work, use the command `yes` at a prompt.

```
/home/larry# yes
```

This will have the startling effect of running a long column of `y`'s down the left hand side of your screen, faster than you can follow.⁵ To get them to stop, you'd normally type `Ctrl-C` to kill it, but instead you should type `Ctrl-Z` this time. It appears to have stopped, but there will be a message before your prompt, looking more or less like this:

⁵There are good reasons for this strange command to exist. Occasional commands ask for confirmation—a "yes" answer to a question. The `yes` command allows a programmer to automate the response to these questions.

```
[1]+ Stopped          yes
```

It means that the process `yes` has been suspended in the background. You can get it running again by typing `fg` at the prompt, which will put it into the foreground again. If you wish, you can do other things first, while it's suspended. Try a few `ls`'s or something before you put it back in the foreground.

Once it's returned to the foreground, the `y`'s will start coming again, as fast as before. You do not need to worry that while you had it suspended it was "storing up" more `y`'s to send to the screen: when a program is suspended the whole program doesn't run until you bring it back to life. (Now type `Ctrl-C` to kill it for good, once you've seen enough).

Let's pick apart that message we got from the shell:

```
[1]+ Stopped          yes
```

The number in brackets is the job number of this job, and will be used when we need to refer to it specifically. (Naturally, since job control is all about running multiple processes, we need some way to tell one from another). The `+` following it tells us that this is the "current job" — that is, the one most recently moved from the foreground to the background. If you were to type `fg`, you would put the job with the `+` in the foreground again. (More on that later, when we discuss running multiple jobs at once). The word `Stopped` means that the job is "stopped". The job isn't dead, but it isn't running right now. LINUX has saved it in a special suspended state, ready to jump back into the action should anyone request it. Finally, the `yes` is the name of the process that has been stopped.

Before we go on, let's kill this job and start it again in a different way. The command is named `kill` and can be used in the following way:

```
/home/larry# kill %1
[1]+ Stopped          yes
/home/larry#
```

That message about it being "stopped" again is misleading. To find out whether it's still alive (that is, either running or frozen in a suspended state), type `jobs`:

```
/home/larry# jobs
[1]+ Terminated      yes
/home/larry#
```

There you have it—the job has been terminated! (It's possible that the `jobs` command showed nothing at all, which just means that there are no jobs running in the background. If you just killed a job, and typing `jobs` shows nothing, then you know the kill was successful. Usually it will tell you the job was "terminated".)

Now, start `yes` running again, like this:

```
/home/larry# yes > /dev/null
```

If you read the section about input and output redirection, you know that this is sending the output of `yes` into the special file `/dev/null`. `/dev/null` is a black hole that eats any output sent to it (you can imagine that stream of `y`'s coming out the back of your computer and drilling a hole in the wall, if that makes you happy).

After typing this, you will not get your prompt back, but you will not see that column of `y`'s either. Although output is being sent into `/dev/null`, the job is still running in the foreground. As usual, you can suspend it by hitting `Ctrl-Z`. Do that now to get the prompt back.

```
/home/larry# yes > /dev/null
["yes" is running, and we just typed ctrl-z]
[1]+ Stopped          yes >/dev/null

/home/larry#
```

Hmm... is there any way to get it to actually run in the background, while still leaving us the prompt for interactive work? The command to do that is `bg`:

```
/home/larry# bg
[1]+ yes >/dev/null &
/home/larry#
```

Now, you'll have to trust me on this one: after you typed `bg`, `yes > /dev/null` began to run again, but this time in the background. In fact, if you do things at the prompt, like `ls` and stuff, you might notice that your machine has been slowed down a little bit (endlessly generating and discarding a steady stream of `y`'s does take some work, after all!) Other than that, however, there are no effects. You can do anything you want at the prompt, and `yes` will happily continue to sending its output into the black hole.

There are now two different ways you can kill it: with the `kill` command you just learned, or by putting the job in the foreground again and hitting it with an interrupt, `Ctrl-C`. Let's try the second way, just to understand the relationship between `fg` and `bg` a little better;

```
/home/larry# fg
yes >/dev/null

[now it's in the foreground again. Imagine that I hit ctrl-c to terminate it]

/home/larry#
```

There, it's gone. Now, start up a few jobs running in simultaneously, like this:

```
/home/larry# yes > /dev/null &
[1] 1024
/home/larry# yes | sort > /dev/null &
[2] 1026
/home/larry# yes | uniq > /dev/null
[and here, type ctrl-z to suspend it, please]
```

```
[3]+ Stopped          yes | uniq >/dev/null
/home/larry#
```

The first thing you might notice about those commands is the trailing `&` at the end of the first two. Putting an `&` after a command tells the shell to start running in the background right from the very beginning. (It's just a way to avoid having to start the program, type `[ctrl-z]`, and then type `bg`.) So, we started those two commands running in the background. The third is suspended and inactive at the moment. You may notice that the machine has become slower now, as the two running ones require some amount of CPU time.

Each one told you it's job number. The first two also showed you their process identification numbers, or PID's, immediately following the job number. The PID's are normally not something you need to know, but occasionally come in handy.

Let's kill the second one, since I think it's making your machine slow. You could just type `kill %2`, but that would be too easy. Instead, do this:

```
/home/larry# fg %2
yes | sort >/dev/null
[type ctrl-c to kill it]

/home/larry#
```

As this demonstrates, `fg` takes parameters beginning with `%` as well. In fact, you could just have typed this:

```
/home/larry# %2
yes | sort >/dev/null
[type ctrl-c to kill it]

/home/larry#
```

This works because the shell automatically interprets a job number as a request to put that job in the foreground. It can tell job numbers from other numbers by the preceding `%`. Now type `jobs` to see which jobs are left running:

```
/home/larry# jobs
[1]-  Running          yes >/dev/null &
[3]+  Stopped         yes | uniq >/dev/null
/home/larry#
```

The `"-"` means that job number 1 is second in line to be put in the foreground, if you just type `fg` without giving it any parameters. The `"+"` means the specified job is first in line—a `fg` without parameters will bring job number 3 to the foreground. However, you can get to it by naming it, if you wish:

```
/home/larry# fg %1
yes >/dev/null
[now type ctrl-z to suspend it]
```

```
[1]+ Stopped          yes >/dev/null
/home/larry#
```

Having changed to job number 1 and then suspending it has also changed the priorities of all your jobs. You can see this with the `jobs` command:

```
/home/larry# jobs
[1]+ Stopped          yes >/dev/null
[3]- Stopped         yes | uniq >/dev/null
/home/larry#
```

Now they are both stopped (because both were suspended with `[ctrl-z]`), and number 1 is next in line to come to the foreground by default. This is because you put it in the foreground manually, and then suspended it. The `"+"` always refers to the most recent job that was suspended from the foreground. You can start it running again:

```
/home/larry# bg
[1]+ yes >/dev/null &
/home/larry# jobs
[1]-  Running          yes >/dev/null
[3]+  Stopped         yes | uniq >/dev/null
/home/larry#
```

Notice that now it is running, and the other job has moved back up in line and has the `+`. Now let's kill them all so your system isn't permanently slowed by processes doing nothing.

```
/home/larry# kill %1 %3
[3]  Terminated      yes | uniq >/dev/null
/home/larry# jobs
[1]+  Terminated     yes >/dev/null
/home/larry#
```

You should see various messages about termination of jobs—nothing dies quietly, it seems. Figure 6.1 on the facing page shows a quick summary of what you should know for job control.

6.4.2 The Theory of Job Control

It is important to understand that job control is done by the shell. There is no program on the system called `fg`; rather, `fg`, `bg`, `&`, `jobs`, and `kill` are all shell-builtins (actually, sometimes `kill` is an independent program, but the `bash` shell used by Linux has it built in). This is a logical way to do it: since each user wants their own job control space, and each user already has their own shell, it is easiest to just have the shell keep track of the user's jobs. Therefore, each user's job numbers are meaningful only to that user: my job number [1] and your job number [1] are probably two totally different processes. In fact, if you are logged in more than once, each of your shells will have unique job control data, so you as a user might have two different jobs with the same number running in two different shells.

Figure 6.1 A summary of commands and keys used in job control.

<code>fg %job</code>	This is a shell command that returns a job to the foreground. To find out which one this is by default, type <code>jobs</code> and look for the one with the <code>+</code> . Parameters: Optional job number. The default is the process identified with <code>+</code> .
<code>&</code>	When an <code>&</code> is added to the end of the command line, it tells the command to run in the background automatically. This process is then subject to all the usual methods of job control detailed here.
<code>bg %job</code>	This is a shell command that causes a suspended job to run in the background. To find out which one this is by default, type <code>jobs</code> and look for the one with the <code>+</code> . Parameters: Optional job number. The default is the process identified with <code>+</code> .
<code>kill %job PID</code>	This is a shell command that causes a background job, either suspended or running, to terminate. You should always specify the job number or PID, and if you are using job numbers, remember to precede them with a <code>%</code> . Parameters: Either the job number (preceded by <code>%</code>) or PID (no <code>%</code> is necessary). More than one process or job can be specified on one line.
<code>jobs</code>	This shell command just lists information about the jobs currently running or suspending. Sometimes it also tells you about ones that have just exited or been terminated.
<code>ctrl-c</code>	This is the generic interrupt character. Usually, if you type it while a program is running in the foreground, it will kill the program (sometimes it takes a few tries). However, not all programs will respond to this method of termination.
<code>ctrl-z</code>	This key combination usually causes a program to suspend, although a few programs ignore it. Once suspended, the job can be run in the background or killed.

The way to tell for sure is to use the Process ID numbers (PID's). These are system-wide — each process has its own unique PID number. Two different users can refer to a process by its PID and know that they are talking about the same process (assuming that they are logged into the same machine!)

Let's take a look at one more command to understand what PIDs are. The `ps` command will list all running processes, including your shell. Try it out. It also has a few options, the most important of which (to many people) are `a`, `u`, and `x`. The `a` option will list processes belonging to any user, not just your own. The `x` switch will list processes that don't have a terminal associated with them.⁶ Finally, the `u` switch will give additionally information about the process that is frequently useful.

To really get an idea of what your system is doing, put them all together: `ps -aux`. You can then see the process that uses the more memory by looking at the `%MEM` column, and the most CPU by looking at the `%CPU` column. (The `TIME` column lists the *total* amount of CPU time used.)

⁶This only makes sense for certain system programs that don't have to talk to users through a keyboard.

Another quick note about PIDs. `kill`, in addition to taking options of the form `%job#`, will take options of raw PIDs. So, put a `yes > /dev/null` in the background, run `ps`, and look for `yes`. Then type `kill PID`.⁷

If you start to program in C on your Linux system, you will soon learn that the shell's job control is just an interactive version of the function calls `fork` and `exec1`. This is too complex to go into here, but may be helpful to remember later on when you are programming and want to run multiple processes from a single program.

6.5 Virtual Consoles: Being in Many Places at Once

Linux supports **virtual consoles**. These are a way of making your single machine seem like multiple terminals, all connected to one Linux kernel. Thankfully, using virtual consoles is one of the simplest things about Linux: there are "hot keys" for switching among the consoles quickly. To try it, log in to your Linux system, hold down the left `Alt` key, and press `F2` (that is, the function key number 2).⁸

You should find yourself at another login prompt. Don't panic: you are now on virtual console (VC) number 2! Log in here and do some things -- a few `ls`'s or whatever -- to confirm that this is a real login shell. Now you can return to VC number 1, by holding down the left `Alt` and pressing `F1`. Or you can move on to a *third* VC, in the obvious way (`Alt-F3`).

Linux systems generally come with four VC's enabled by default. You can increase this all the way to eight; this should be covered in *The Linux System Administrator's Guide*. It involves editing a file in `/etc` or two. However, four should be enough for most people.

Once you get used to them, VC's will probably become an indispensable tool for getting many things done at once. For example, I typically run Emacs on VC 1 (and do most of my work there), while having a communications program up on VC 3 (so I can be downloading or uploading files by modem while I work, or running jobs on remote machines), and keep a shell up on VC 2 just in case I want to run something else without tying up VC 1.

⁷In general, it's easier to just kill the job number instead of using PIDs.

⁸Make sure you are doing this from text consoles: if you are running X windows or some other graphical application, it probably won't work, although rumor has it that X Windows will soon allow virtual console switching under Linux.

Chapter 7

Powerful Little Programs

```
better !pout !cry
better watchout
lpr why
santa claus <north pole >town

cat /etc/passwd >list
ncheck list
ncheck list
cat list | grep naughty >nogiftlist
cat list | grep nice >giftlist
santa claus <north pole > town

who | grep sleeping
who | grep awake
who | egrep 'bad|good'
for (goodness sake) (
    be good
)
```

7.1 The Power of Unix

The power of Unix is hidden in small commands that don't seem too useful when used alone, but when combined with other commands (either directly or indirectly) produce a system that's much more powerful and flexible than most other operating systems. The commands I'm going to talk about in this chapter include `sort`, `grep`, `more`, `cat`, `wc`, `spell`, `diff`, `head`, and `tail`. Unfortunately, it isn't totally intuitive what these names mean right now.

Let's cover what each of these utilities do separately and then I'll give some examples of how to use them together.¹

¹Please note that the short summaries on commands in this chapter are not comprehensive. Please consult the

7.2 Operating on Files

In addition to the commands like `cd`, `mv`, and `rm` you learned in Chapter 4, there are other commands that just operate on files but not the data in them. These include `touch`, `chmod`, `du`, and `df`. All of these files don't care what is in the file--the merely change some of the things Unix remembers about the file.

Some of the things these commands manipulate:

- The time stamp. Each file has three dates associated with it.² The three dates are the creation time (when the file was created), the last modification time (when the file was last changed), and the last access time (when the file was last read).
- The owner. Every file in Unix is owned by one user or the other.
- The group. Every file also has a group of users it is associated with. The most common group for user files is called `users`, which is usually shared by all the user account on the system.
- The permissions. Every file has permissions (sometimes called "privileges") associated with it which tell Unix who can access what file, or change it, or, in the case of programs, execute it. Each of these permissions can be toggled separately for the owner, the group, and all other users.

`touch file1 file2 ... fileN`

`touch` will update the time stamps of the files listed on the command line to the current time. If a file doesn't exist, `touch` will create it. It is also possible to specify the time that `touch` will set files to--consult the the manpage for `touch`.

`chmod [-Rfv] mode file1 file2 ... fileN`

The command used to change the permissions on a file is called `chmod`, short for `change mode`. Before I go into how to use the command, let's discuss what permissions are in Unix. Each file has a group of permissions associated with it. These permissions tell Unix whether or not the file can be read from, written to, or executed as a program. (In the next few paragraphs, I'll talk about users doing these things. Any programs a user runs are allowed to do the same things a user is. This can be a security problem if you don't know what a particular program does.)

Unix recognizes three different types of people: first, the owner of the file (and the person allowed to use `chmod` on that file). Second, the "group". The group of most of your files might be "users", meaning the normal users of the system. (To find out the group of a particular file, use `ls -l file`.)

command's manpage if you want to know every option.

²Older filesystems in LINUX only stored one date, since they were derived from Minix. If you have one of these filesystems, some of the information will merely be unavailable--operation will be mostly unchanged.

Then, there's everybody else who isn't the owner and isn't a member of the group, appropriately called "other".

So, a file could have read and write permissions for the owner, read permissions for the group, and no permissions for all others. Or, for some reason, a file could have read/write permissions for the group and others, but *no* permissions for the owner!

Let's try using `chmod` to change a few permissions. First, create a new file using `cat`, `emacs`, or any other program. By default, you'll be able to read and write this file. (The permissions given other people will vary depending on how the system and your account is setup.) Make sure you can read the file using `cat`. Now, let's take away your read privilege by using `chmod u-r filename`. (The parameter `u-r` decodes to "user minus read".) Now if you try to read the file, you get a `Permission denied` error! Add read privileges back by using `chmod u+r filename`.

Directory permissions use the same three ideas: read, write, and execute, but act slightly differently. The read privilege allows the user (or group or others) to read the directory list the names of the files. The write permission allows the user (or group or others) to add or remove files. The execute permission allows the user to access files in the directory or any subdirectories. (If a user doesn't have execute permissions for a directory, they can't even `cd` to it!)

To use `chmod`, replace the *mode* with what to operate on, either user, group, other, or all, and what to do with them. (That is, use a plus sign to indicate adding a privilege or a minus sign to indicate taking one away. Or, an equals sign will specify the exact permissions.) The possible permissions to add are read, write, and execute.

`chmod's` `R` flag will change a directory's permissions, and all files in that directory, and all subdirectories, all the way down the line. (The '`R`' stands for recursive.) The `f` flag forces `chmod` to attempt to change permissions, even if the user isn't the owner of the file. (If `chmod` is given the `f` flag, it won't print an error message when it fails to change a file's permissions.) The `v` flag makes `chmod` verbose—it will report on what it's done.

7.3 System Statistics

Commands in this section will display statistics about the operating system, or a part of the operating system.

`du [-abs] [path1 path2 ... pathN]`

`du` stands for disk usage. It will count the amount of disk space a given directory and all its subdirectories take up on the disk. `du` by itself will return a list of how much space every subdirectory of the current directory consumes, and, at the very bottom, how much space the current directory (plus all the previously counted subdirectories) use. If you give it a parameter or two, it will count the amount of space used by those files or directories instead of the current one.

The `a` flag will display a count for files, as well as directories. An option of `b` will display, instead of kilobytes (1024 characters), the total in bytes. One byte is the equivalent of one letter in a text

document. And the `s` flag will just display the directories mentioned on the command-line and *not* their subdirectories.

`df`

`df` is short for "disk filling": it summarizes the amount of disk space in use. For each filesystem (remember, different filesystems are either on different drives or partitions) it shows the total amount of disk space, the amount used, the amount available, and the total capacity of the filesystem that's used.

One odd thing you might encounter is that it's possible for the capacity to go over 100%, or the used plus the available not to equal the total. This is because Unix reserves some space on each filesystem for root. That way if a user accidentally fills the disk, the system will still have a little room to keep on operating.

For most people, `df` doesn't have any useful options.

`uptime`

The `uptime` program does exactly what one would suspect. It prints the amount of time the system has been "up"—the amount of time from the last Unix boot.

`uptime` also gives the current time and the load average. The load average is the average number of jobs waiting to run in a certain time period. `uptime` displays the load average for the last minute, five minutes, and ten minutes. A load average near zero indicates the system has been relatively idle. A load average near one indicates that the system has been almost fully utilized but nowhere near overtaxed. High load averages are the result of several programs being run simultaneously.

Amazingly, `uptime` is one of the few Unix programs that have *no* options!

`who`

`who` displays the current users of the system and when they logged in. If given the parameters `am i` (as in: `who am i`), it displays the current user.

`w [-f] [username]`

The `w` program displays the current users of the system and what they're doing. (It basically combines the functionality of `uptime` and `who`. The header of `w` is exactly the same as `uptime`, and each line shows a user, when the logged on (and how long they've been idle). `JCPU` is the total amount of CPU time used by that user, while `PCPU` the total amount of CPU time used by their present task.

If `v` is given the option `f`, it shows the remote system they logged in from, if any. The optional parameter restricts `v` to showing only the named user.

7.4 What's in the File?

There are two major commands used in Unix for listing files, `cat` and `more`. I've talked about both of them in Chapter 6.

```
cat [-nA] [file1 file2 ... fileN]
```

`cat` is not a user friendly command—it doesn't wait for you to read the file, and is mostly used in conjunction with pipes. However, `cat` does have some useful command-line options. For instance, `n` will number all the lines in the file, and `A` will show control characters as normal characters instead of (possibly) doing strange things to your screen. (Remember, to see some of the stranger and perhaps "less useful" options, use the `man` command: `man cat`.) `cat` will accept input from `stdin` if no files are specified on the command-line.

```
more [-l] [+linenumber] [file1 file2 ... fileN]
```

`more` is much more useful, and is the command that you'll want to use when browsing ASCII text files. The only interesting option is `l`, which will tell `more` that you aren't interested in treating the character `Ctrl-L` as a "new page" character. `more` will start on a specified linenumber.

Since `more` is an interactive command, I've summarized the major interactive commands below:

`Spacebar` Moves to the next screen of text.

`d` This will scroll the screen by 11 lines, or about half a normal, 25-line, screen.

`/` Searches for a regular expression. While a regular expression can be quite complicated, you can just type in a text string to search for. For example, `/toad``return` would search for the next occurrence of "toad" in your current file. A slash followed by a return will search for the next occurrence of what you last searched for.

`n` This will also search for the next occurrence of your regular expression.

`:n` If you specified more than one file on the command line, this will move to the next file.

`:p` This will move to the previous file.

`q` Exits from `more`.

```
head [-lines] [file1 file2 ... fileN]
```

`head` will display the first ten lines in the listed files, or the first ten lines of `stdin` if no files are specified on the command line. Any numeric option will be taken as the number of lines to print, so `head -15 frog` will print the first fifteen lines of the file `frog`.

```
tail [-lines] [file1 file2 ... fileN]
```

Like `head`, `tail` will display only a fraction of the file. Naturally, `tail` will display the end of the file, or the last ten lines that come through `stdin`. `tail` also accepts a option specifying the number of lines.

```
file [file1 file2 ... fileN]
```

The `file` command attempts to identify what format a particular file is written in. Since not all files have extensions or other easy to identify marks, the `file` command performs some rudimentary checks to try and figure out exactly what it contains.

Be careful, though, because it is quite possible for `file` to make a wrong identification.

7.5 Information Commands

This section discusses the commands that will alter a file, perform a certain operation on the file, or display statistics on the file.

```
grep [-nvwx] [-number] expression [file1 file2 ... fileN]
```

One of the most useful commands in Unix is `grep`, the generalized regular expression parser. This is a fancy name for a utility which can only search a text file. The easiest way to use `grep` is like this:

```
/home/larry# cat animals
Animals are very interesting creatures. One of my favorite animals is
the tiger, a fearsome beast with large teeth.
I also like the lion---it's really neat!
/home/larry# grep iger animals
the tiger, a fearsome beast with large teeth.
/home/larry#
```

One disadvantage of this is, although it shows you all the lines containing your word, it doesn't

tell you where to look in the file—no line number. Depending on what you're doing, this might be fine. For instance, if you're looking for errors from a program's output, you might try `a.out | grep error`, where `a.out` is your program's name.

If you're interested in where the match(es) are, use the `n` switch to `grep` to tell it to print line numbers. Use the `v` switch if you want to see all the lines that *don't* match the specified expression.

Another feature of `grep` is that it matches only parts of a word, like my example above where `iger` matched `tiger`. To tell `grep` to only match whole words, use the `w`, and the `x` switch will tell `grep` to only match whole lines.

If you don't specify any files, `grep` will examine `stdin`.

```
wc [-clw] [file1 file2 ... fileN]
```

`wc` stands for word count. It simply counts the number of words, lines, and characters in the file(s). If there aren't any files specified on the command line, it operates on `stdin`.

The three parameters, `clw`, stand for character, line, and word respectively, and tell `wc` which of the three to count. Thus, `wc -cw` will count the number of characters and words, but not the number of lines. `wc` defaults to counting everything—words, lines, and characters.

One nice use of `wc` is to find how many files are in the present directory: `ls | wc -w`. If you wanted to see how many files that ended with `.c` there are, try `ls *.c | wc -w`.

```
spell [file1 file2 ... fileN]
```

`spell` is a very simple Unix spelling program, usually for American English.³ `spell` is a filter, like most of the other programs we've talked about, which sucks in an ASCII text file and outputs all the words it considers misspellings. `spell` operates on the files listed in the command line, or, if there weren't any there, `stdin`.

A more sophisticated spelling program, `ispell` is probably also available on your machine. `ispell` will offer possible correct spellings and a fancy menu interface if a filename is specified on the command line or will run as a filter-like program if no files are specified.

While operation of `ispell` should be fairly obvious, consult the man page if you need more help.

```
cmp file1 [file2]
```

`cmp` compares two files. The first must be listed on the command line, while the second is either listed as the second parameter or is read in from standard input. `cmp` is very simple, and merely tells you where the two files first differ.

³While there are versions of this for several other European languages, the copy on your LINUX machine is most likely for American English.

```
diff file1 file2
```

One of the most complicated standard Unix commands is called `diff`. The GNU version of `diff` has over twenty command line options! It is a much more powerful version of `cmp` and shows you what the differences are instead of merely telling you where the first one is.

Since talking about even a good portion of `diff` is beyond the scope of this book, I'll just talk about the basic operation of `diff`. In short, `diff` takes two parameters and displays the differences between them on a line-by-line basis. For instance:

```
/home/larry# cat frog
Animals are very interesting creatures. One of my favorite animals is
the tiger, a fearsome beast with large teeth.
I also like the lion---it's really neat!
/home/larry# cp frog toad
/home/larry# diff frog toad
/home/larry# cat dog
Animals are very interesting creatures. One of my favorite animals is

the tiger, a fearsome beast with large teeth.
I also like the lion---it's really neat!
/home/larry# diff frog dog
1c1,2
< Animals are very interesting creatures. One of my favorite animals is
---
> Animals are very interesting creatures. One of my favorite animals is
>
3c4
< I also like the lion---it's really neat!
---
> I also like the lion---it's really neat!
/home/larry#
```

As you can see, `diff` outputs nothing when the two files are identical. Then, when I compared two different files, it had a section header, `1c1,2` saying it was comparing line 1 of the left file, `frog`, to lines 1-2 of `dog` and what differences it noticed. Then it compared line 3 of `frog` to line 4 of `dog`. While it may seem strange at first to compare different line numbers, it is much more efficient than listing out every single line if there is an extra return early in one file.

```
gzip [-v#] [file1 file2 ... fileN]
gunzip [-v] [file1 file2 ... fileN]
zcat [file1 file2 ... fileN]
```

These three programs are used to compress and decompress data. `gzip`, or GNU Zip, is the

program that reads in the original file(s) and outputs files that are smaller. `gzip` deletes the files specified on the command line and replaces them with files that have an identical name except that they have “.gz” appended to them.

tr *string1 string2*

The “translate characters” command operates on standard input—it doesn’t accept a filename as a parameter. Instead, its two parameters are arbitrary strings. It replaces all occurrences of *string1* in the input with *string2*. In addition to relatively simple commands such as `tr frog toad`, `tr` can accept more complicated commands. For instance, here’s a quick way of converting lowercase characters into uppercase ones:

```
/home/larry# tr [:lower:] [:upper:]  
this is a WEIRD sentence.  
THIS IS A WEIRD SENTENCE.
```

`tr` is fairly complex and usually used in small shell programs.

Chapter 9

I Gotta Be Me!

If God had known we'd need foresight, she would have given it to us.

9.1 bash Customization

One of the distinguishing things about the Unix philosophy is that the system's designers did not attempt to predict every need that users might have; instead, they tried to make it easy for each individual user to tailor the environment to their own particular needs. This is mainly done through **configuration files**. These are also known as "init files", "rc files" (for "run control"), or even "dot files", because the filenames often begin with ".". If you'll recall, filenames that start with "." aren't normally displayed by `ls`.

The most important configuration files are the ones used by the shell. Linux's default shell is `bash`, and that's the shell this chapter covers. Before we go into how to customize `bash`, we should know what files `bash` looks at.

9.1.1 Shell Startup

There are several different ways `bash` can run. It can run as a **login shell**, which is how it runs when you first login. The login shell should be the first shell you see.

Another way `bash` can run is as an **interactive shell**. This is any shell which presents a prompt to a human and waits for input. A login shell is also an interactive shell. A way you can get a non-login interactive shell is, say, a shell inside `xterm`. Any shell that was created by some other way besides logging in is a non-login shell.

Finally, there are **non-interactive shells**. These shells are used for executing a file of commands, much like MS-DOS's batch files—the files that end in `.BAT`. These **shell scripts** function like mini-programs. While they are usually much slower than a regular compiled program, it is often true that they're easier to write.

Depending on the type of shell, different files will be used at shell startup:

Type of Shell	Action
Interactive login	The file <code>.bash_profile</code> is read and executed
Interactive	The file <code>.bashrc</code> is read and executed
Non-interactive	The shell script is read and executed

9.1.2 Startup Files

Since most users want to have largely the same environment no matter what type of interactive shell they wind up with, whether or not it's a login shell, we'll start our configuration by putting a very simple command into our `.bash_profile`: `source ~/.bashrc`. The `source` command tells the shell to interpret the argument as a shell script. What it means for us is that everytime `.bash_profile` is run, `.bashrc` is *also* run.

Now, we'll just add commands to our `.bashrc`. If you ever want a command to only be run when you login, add it to your `.bash_profile`.

9.1.3 Aliasing

What are some of the things you might want to customize? Here's something that I think about 90% of Dash users have put in their `.bashrc`:

```
alias ll="ls -l"
```

That command defined a shell alias called `ll` that "expands" to the normal shell command `"ls -l"` when invoked by the user. So, assuming that `Dash` has read that command in from your `.bashrc`, you can just type `ll` to get the effect of `"ls -l"` in only half the keystrokes. What happens is that when you type `ll` and hit **Return**, `Dash` intercepts it, because it's watching for aliases, replaces it with `"ls -l"`, and runs that instead. There is no actual program called `ll` on the system, but the shell automatically translated the alias into a valid program.

Some sample aliases are in Figure 9.1.3. You could put them in your own `.bashrc`. One especially interesting alias is the first one. With that alias, whenever someone types `ls`, they automatically have a `-F` flag tacked on. (The alias doesn't try to expand itself again.) This is a common way of adding options that you use every time you call a program.

Notice the comments with the `#` character in Figure 9.1.3. Whenever a `#` appears, the shell ignores the rest of the line.

You might have noticed a few odd things about them. First of all, I leave off the quotes in a few of the aliases—like `pu`. Strictly speaking, quotes aren't necessary when you only have one word on the right of the equal sign.

It never hurts to have quotes either, so don't let me get you into any bad habits. You should certainly use them if you're going to be aliasing a command with options and/or arguments:

```
alias rf="refrobnicate -verbose -prolix -wordy -o foo.out"
```

Figure 9.1 Some sample aliases for bash.

```
alias ls="ls -F"          # give characters at the end of listing
alias ll="ls -l"         # special ls
alias la="ls -a"
alias ro="rm *"; rm .*"  # this removes backup files created by Emacs
alias rd="rmdir"         # saves typing!
alias md="mkdir"
alias pu=pushd           # pushd, popd, and dirs weren't covered in this
alias po=popd            # manual--you might want to look them up
alias ds=dirs            # in the bash manpage
# these all are just keyboard shortcuts
alias to="telnet cs.oberlin.edu"
alias ta="telnet altair.mcs.anl.gov"
alias tg="telnet wombat.gnu.ai.mit.edu"
alias tko="tpalk kold@cs.oberlin.edu"
alias tjo="talk jim@cs.oberlin.edu"
alias mroe="more"        # spelling correction!
alias moer="more"
alias email="emacs -f rmail" # my mail reader
alias ed2="emacs -d floss:0 -fg \"grey95\" -bg \"grey50\""
# one way of invoking emacs
```

Also, the final alias has some funky quoting going on:

```
alias ed2="emacs -d floss:0 -fg \"grey95\" -bg \"grey50\""
```

As you might have guessed, I wanted to pass double-quotes in the options themselves, so I had to quote those with a backslash to prevent bash from thinking that they signaled the end of the alias.

Finally, I have actually aliased two common typing mistakes, "mroe" and "moer", to the command I meant to type, more. Aliases do not interfere with your passing arguments to a program. The following works just fine:

```
/home/larry# mroe hurd.txt
```

In fact, knowing how to make your own aliases is probably at least half of all the shell customization you'll ever do. Experiment a little, find out what long commands you find yourself typing frequently, and make aliases for them. You'll find that it makes working at a shell prompt a much more pleasant experience.

9.1.4 Environment Variables

Another major thing one does in a .bashrc is set environment variables. And what are environment variables? Let's go at it from the other direction: suppose you are reading the documentation for the program fruggle, and you run across these sentences:

Fruggle normally looks for its configuration file, .frugglerc, in the user's home directory. However, if the environment variable FRUGGLEPATH is set to a different filename, it will look there instead.

Every program executes in an environment, and that environment is defined by the shell that called the program¹. The environment could be said to exist "within" the shell. Programmers have a special routine for querying the environment, and the fruggle program makes use of this routine. It checks the value of the environment variable FRUGGLEPATH. If that variable turns out to be undefined, then it will just use the file .frugglerc in your home directory. If it is defined, however, fruggle will use the variable's value (which should be the name of a file that fruggle can use) instead of the default .frugglerc.

Here's how you can change your environment in bash:

```
/home/larry# export PGPPATH=/home/larry/secrets/pgp
```

You may think of the export command as meaning "Please export this variable out to the environment where I will be calling programs, so that its value is visible to them." There are actually reasons to call it export, as you'll see later.

This particular variable is used by Phil Zimmerman's infamous public-key encryption program, pgp. By default, pgp uses your home directory as a place to find certain files that it needs (containing encryption keys), and also as a place to store temporary files that it creates when it's running. By setting variable PGPPATH to this value, I have told it to use the directory /home/larry/secrets/pgp instead. I had to read the pgp manual to find out the exact name of the variable and what it does, but it is fairly standard to use the name of the program in capital letters, prepended to the suffix "PATH".

It is also useful to be able to query the environment:

```
/home/larry# echo $PGPPATH
/home/larry/.pgp
/home/larry#
```

Notice the "\$": you prefix an environment variable with a dollar sign in order to extract the variable's value. Had you typed it without the dollar sign, echo would have simply echoed its argument(s):

```
/home/larry# echo PGPPATH
PGPPATH
/home/larry#
```

The "\$" is used to evaluate environment variables, but it only does so in the context of the shell-- that is, when the shell is interpreting. When is the shell interpreting? Well, when you are

¹Now you see why shells are so important. Imagine if you had to pass a whole environment by hand every time you called a program!

Figure 9.2 Some important environment variables.

Variable name	Contains	Example
HOME	Your home directory	/home/larry
TERM	Your terminal type	xterm, vt100, or console
SHELL	The path to your shell	/bin/bash
USER	Your login name	larry
PATH	A list to search for programs	/bin:/usr/bin:/usr/local/bin:/usr/bin/X11

typing commands at the prompt, or when `bash` is reading commands from a file like `.bashrc`, it can be said to be “interpreting” the commands.

There's another command that's very useful for querying the environment: `env`. `env` will merely list all the environment variables. It's possible, especially if you're using `X`, that the list will scroll off the screen. If that happens, just pipe `env` through `more`: `env | more`.

A few of these variables can be fairly useful, so I'll cover them. Look at Figure 9.1.4. Those four variables are defined automatically when you login: you don't set them in your `.bashrc` or `.bash_login`.

Let's take a closer look at the `TERM` variable. To understand that one, let's look back into the history of Unix: The operating system needs to know certain facts about your console, in order to perform basic functions like writing a character to the screen, moving the cursor to the next line, etc. In the early days of computing, manufacturers were constantly adding new features to their terminals: first reverse-video, then maybe European character sets, eventually even primitive drawing functions (remember, these were the days before windowing systems and mice). However, all of these new functions represented a problem to programmers: how could they know what a terminal supported and didn't support? And how could they support new features without making old terminals worthless?

In Unix, the answer to these questions was `/etc/termcap`. `/etc/termcap` is a list of all of the terminals that your system knows about, and how they control the cursor. If a system administrator got a new terminal, all they'd have to do is add an entry for that terminal into `/etc/termcap` instead of rebuilding all of Unix. Sometimes, it's even simpler. Along the way, Digital Equipment Corporation's vt100 terminal became a pseudo-standard, and many new terminals were built so that they could emulate it, or behave as if they were a vt100.

Under LINUX, `TERM`'s value is sometimes `console`, which is a vt100-like terminal with some extra features.

Another variable, `PATH`, is also crucial to the proper functioning of the shell. Here's mine:

```
/home/larry# env | grep ^PATH
PATH=/home/larry/bin:/bin:/usr/bin:/usr/local/bin:/usr/bin/X11:/usr/TeX/bin
/home/larry#
```

Your `PATH` is a colon-separated list of the directories the shell should search for programs, when you type the name of a program to run. When I type `ls` and hit `[Return]`, for example, the Bash

first looks in `/home/larry/bin`, a directory I made for storing programs that I wrote. However, I didn't write `ls` (in fact, I think it might have been written before I was born!). Failing to find it in `/home/larry/bin`, Bash looks next in `/bin`—and there it has a hit! `/bin/ls` does exist and is executable, so Bash stops searching for a program named `ls` and runs it. There might well have been another `ls` sitting in the directory `/usr/bin`, but `bash` would never run it unless I asked for it by specifying an explicit pathname:

```
/home/larry# /usr/bin/ls
```

The `PATH` variable exists so that we don't have to type in complete pathnames for every command. When you type a command, Bash looks for it in the directories named in `PATH`, in order, and runs it if it finds it. If it doesn't find it, you get a rude error:

```
/home/larry# clubly
clubly: command not found
```

Notice that my `PATH` does not have the current directory, `“.”`, in it. If it did, it might look like this:

```
/home/larry# echo $PATH
.: /home/larry/bin:/bin:/usr/bin:/usr/local/bin:/usr/bin/X11:/usr/TeX/bin
/home/larry#
```

This is a matter of some debate in Unix-circles (which you are now a member of, whether you like it or not). The problem is that having the current directory in your path can be a security hole. Suppose that you `cd` into a directory where somebody has left a “Trojan Horse” program called `ls`, and you do an `ls`, as would be natural on entering a new directory. Since the current directory, `“.”`, came first in your `PATH`, the shell would have found this version of `ls` and executed it. Whatever mischief they might have put into that program, you have just gone ahead and executed (and that could be quite a lot of mischief indeed). The person did not need root privileges to do this; they only needed write permission on the directory where the “false” `ls` was located. It might even have been their home directory, if they knew that you would be poking around in there at some point.

On your own system, it's highly unlikely that people are leaving traps for each other. All the users are probably friends or colleagues of yours. However, on a large multi-user system (like many university computers), there could be plenty of unfriendly programmers whom you've never met. Whether or not you want to take your chances by having `“.”` in your path depends on your situation; I'm not going to be dogmatic about it either way, I just want you to be aware of the risks involved². Multi-user systems really are communities, where people can do things to one another in all sorts of unforeseen ways.

The actual way that I set my `PATH` involves most of what you've learned so far about environment variables. Here is what is actually in my `.bashrc`:

```
export PATH=${PATH}.:${HOME}/bin:/bin:/usr/bin:/usr/local/bin:/usr/bin/X11:/usr/TeX/bin
```

²Remember that you can always execute programs in the current directory by being explicit about it, i.e.: `“./foo”`.

Here, I am taking advantage of the fact that the `HOME` variable is set before Bash reads my `.bashrc`, by using its value in setting my `PATH`. The curly braces (“{...}”) are a further level of quoting; they delimit the extent of what the “\$” is to evaluate, so that the shell doesn’t get confused by the text immediately following it (“/bin” in this case). Here is another example of the effect they have:

```
/home/larry# echo ${HOME}foo
/home/larryfoo
/home/larry#
```

Without the curly braces, I would get nothing, since there is no environment variables named `HOMEfoo`.

```
/home/larry# echo $HOMEfoo
/home/larry#
```

Let me clear one other thing up in that path: the meaning of “\$PATH”. What that does is includes the value of any `PATH` variable *previously* set in my new `PATH`. Where would the old variable be set? The file `/etc/profile` serves as a kind of global `.bash_profile` that is common to all users. Having one centralized file like that makes it easier for the system administrator to add a new directory to everyone’s `PATH` or something, without them all having to do it individually. If you include the old path in your new path, you won’t lose any directories that the system already setup for you.

You can also control what your prompt looks like. This is done by setting the value of the environment variable `PS1`. Personally, I want a prompt that shows me the path to the current working directory—here’s how I do it in my `.bashrc`:

```
export PS1='`PWD` '
```

As you can see, there are actually *two* variables being used here. The one being set is `PS1`, and it is being set to the value of `PWD`, which can be thought of as either “Print Working Directory” or “Path to Working Directory”. But the evaluation of `PWD` takes place inside single quotes. The single quotes serve to evaluate the expression inside them, which itself evaluates the variable `PWD`. If you just did `export PS1=$PWD`, your prompt would constantly display the path to the current directory *at the time that PS1 was set*, instead of constantly updating it as you change directories. Well, that’s sort of confusing, and not really all that important. Just keep in mind that you need the quotes if you want the current directory displayed in your prompt.

You might prefer `export PS1='`PWD`>'`, or even the name of your system: `export PS1='hostname`>'`. Let me dissect that last example a little further.

That last example used a *new* type of quoting, the back quotes. These don’t protect something — in fact, you’ll notice that “hostname” doesn’t appear anywhere in the prompt when you run that. What actually happens is that the command inside the backquotes gets evaluated, and the output is put in place of the backquotes and the command name.

Try `echo `ls`` or `wc `ls``. As you get more experienced using the shell, this technique gets more and more powerful.

There’s a lot more to configuring your `.bashrc`, and not enough room to explain it here. You can read the bash man page for more, or ask questions of experienced Bash users. Here is a complete `.bashrc` for you to study; it’s fairly standard, although the search path is a little long.

```
# some random stuff:
ulimit -c unlimited
export history_control=ignoredups
export PS1='`PWD`>'
umask 022

# application-specific paths:
export MANPATH=/usr/local/man:/usr/man
export INFOPATH=/usr/local/info
export GPPATH=${HOME}/.pgp

# make the main PATH:
homepath=${HOME}/~bin
stdpath=/bin:/usr/bin:/usr/local/bin:/usr/ucb:/etc:/usr/etc:/usr/games
pubpath=/usr/public/bin:/usr/gnuoft/bin:/usr/local/contribs/bin
softpath=/usr/bin/X11:/usr/local/bin/X11:/usr/Tex/bin
export PATH=:${homepath}:${stdpath}:${pubpath}:${softpath}
# Technically, the curly braces were not necessary, because the colons
# were valid delimiters; nevertheless, the curly braces are a good
# habit to get into, and they can't hurt.

# aliases
alias ls="ls -CF"
alias fg1="fg %1"
alias fg2="fg %2"
alias tba="talk sussman@tern.mcs.anl.gov"
alias tko="talk kold@ca.oberlin.edu"
alias tji="talk jim@totoro.bio.indiana.edu"
alias mroe="more"
alias moer="more"
alias email="emacs -f vm"
alias pu=pushd
alias po=popd
alias b="~/b"
alias ds=dirs
alias ro="rm *; rm .*"
alias rd="rmdir"
alias ll="ls -l"
alias la="ls -a"
alias rr="rm -r"
alias md="mkdir"
alias ed2="emacs -d floss:0 -fg \"grey95\" -bg \"grey50\""

function gco
```

```
(
gcc -o $1 $1.c -g
)
```

9.2 The X Window System Init Files



Most people prefer to do their work inside a graphical environment, and for Unix machines, that usually means using X. If you're accustomed to the Macintosh or to Microsoft Windows, the X Window System may take a little getting used to, especially in how it is customized.

With the Macintosh or Microsoft Windows, you customize the environment from *within* the environment: if you want to change your background, for example, you do by clicking on the new color in some special graphical setup program. In X, system defaults are controlled by text files, which you edit directly — in other words, you'd type the actual color name into a file in order to set your background to that color.

There is no denying that this method just isn't as slick as some commercial windowing systems. I think this tendency to remain text-based, even in a graphical environment, has to do with the fact that X was created by a bunch of programmers who simply weren't trying to write software that their grandparents could use. This tendency may change in future versions of X (at least I hope it will), but for now, you just have to learn to deal with more text files. It does at least give you very flexible and precise control over your configuration.

Here are the most important files for configuring X:

```
.xinitrc  A script run by X when it starts up.
.twmrc    Read by an X window manager, twm.
.fvwmrc   Read by an X window manager, fvwm.
```

All of these files should be located in your home directory, if they exist at all.

The `.xinitrc` is a simple shell script that gets run when X is invoked. It can do anything any other shell script can do, but of course it makes the most sense to use it for starting up various X programs and setting window system parameters. The last command in the `.xinitrc` is usually the name of a window manager to run, for example `/usr/bin/X11/twm`.

What sort of thing might you want to put in a `.xinitrc` file? Perhaps some calls to the `xsetroot` program, to make your root (background) window and mouse cursor look the way you want them to look. Calls to `xmodmap`, which tells the server³ how to interpret the signals from your keyboard. Any other programs you want started every time you run X (for example, `xclock`).

Here is some of my `.xinitrc`; yours will almost certainly look different, so this is meant only as an example:

```
#!/bin/sh
# The first line tells the operating system which shell to use in
```

³The "server" just means the main X process on your machine, the one with which all other X programs must communicate in order to use the display. These other programs are known as "clients", and the whole deal is called a "client-server" system.

```
# interpreting this script. The script itself ought to be marked as
# executable; you can make it so with "chmod +x ~/.xinitrc".
```

```
# xmodmap is a program for telling the X server how to interpret your
# keyboard's signals. It is *definitely* worth learning about. You
# can do "man xmodmap", "xmodmap -help", "xmodmap -grammar", and more.
# I don't guarantee that the expressions below will mean anything on
# your system (I don't even guarantee that they mean anything on
# mine):
```

```
xmodmap -e 'clear Lock'
xmodmap -e 'keycode 176 = Control_R'
xmodmap -e 'add control = Control_R'
xmodmap -e 'clear Mod2'
xmodmap -e 'add Mod1 = Alt_L Alt_R'
```

```
# xset is a program for setting some other parameters of the X server:
xset m 3 2 & # mouse parameters
xset s 600 5 & # screen saver prefs
xset s noblank & # ditto
xset fp /home/larry/x/fonts # for cxterm
# To find out more, do "xset -help".
```

```
# Tell the X server to superimpose fish.cursor over fish.mask, and use
# the resulting pattern as my mouse cursor:
xsetroot -cursor /home/lab/larry/x/fish.cursor /home/lab/larry/x/fish.mask &
```

```
# a pleasing background pattern and color:
xsetroot -bitmap /home/lab/larry/x/pyramid.xbm -bg tan
```

```
# todo: xrdb here? What about .xdefaults file?
```

```
# You should do "man xsetroot", or "xsetroot -help" for more
# information on the program used above.
```

```
# A client program, the imposing circular color-clock by Jim Blandy:
/usr/local/bin/circles &
```

```
# Maybe you'd like to have a clock on your screen at all times?
/usr/bin/X11/xclock -digital &
```

```
# Allow client X programs running at occs.cs.oberlin.edu to display
# themselves here, do the same thing for juju.mcs.anl.gov:
xhost occs.cs.oberlin.edu
xhost juju.mcs.anl.gov
```

```
# You could simply tell the X server to allow clients running on any
# other host (a host being a remote machine) to display here, but this
# is a security hole -- those clients might be run by someone else,
```

```

# and watch your keystrokes as you type your password or something!
# However, if you wanted to do it anyway, you could use a "+" to stand
# for all possible hostnames, instead of a specific hostname, like
# this:
# xhost +

# And finally, run the window manager:
/usr/bin/X11/twm
# Some people prefer other window managers. I use twm, but fvwm is
# often distributed with Linux too:
# /usr/bin/X11/fvwm

```

Notice that some commands are run in the background (i.e.: they are followed with a "&"), while others aren't. The distinction is that some programs will start when you start X and keep going until you exit—these get put in the background. Others execute once and then exit immediately. `xsetroot` is one such; it just sets the root window or cursor or whatever, and then exits.

Once the window manager has started, it will read its own init file, which controls things like how your menus are set up, which positions windows are brought up at, icon control, and other earth-shakingly important issues. If you use `twm`, then this file is `.twmrc` in your home directory. If you use `fvwm`, then it's `.fvwmrc`, etc. I'll deal with only those two, since they're the window managers you'll be most likely to encounter with Linux.

9.2.1 Twm Configuration

The `.twmrc` is not a shell script—it's actually written in a language specially made for `twm`, believe it or not!⁴ The main thing people like to play with in their `.twmrc` is window style (colors and such), and making cool menus, so here's an example `.twmrc` that does that:

```

# Set colors for the various parts of windows. This has a great
# impact on the "feel" of your environment.
Color
{
    BorderColor "OrangeRed"
    BorderTileForeground "Black"
    BorderTileBackground "Black"
    TitleForeground "black"
    TitleBackground "gold"
    MenuForeground "black"
    MenuBackground "LightGrey"
    MenuItemForeground "LightGrey"
    MenuItemBackground "LightSlateGrey"
    MenuShadowColor "black"

```

⁴This is one of the harsh facts about init files: they generally each have their own idiosyncratic command language. This means that users get very good at learning command languages quickly. I suppose that it would have been nice if early Unix programmers had agreed on some standard init file format, so that we wouldn't have to learn new syntaxes all the time, but to be fair it's hard to predict what kinds of information programs will need.

```

    IconForeground "DimGray"
    IconBackground "Gold"
    IconBorderColor "OrangeRed"
    IconManagerForeground "black"
    IconManagerBackground "honeydew"
}

# I hope you don't have a monochrome system, but if you do...
Monochrome
{
    BorderColor "black"
    BorderTileForeground "black"
    BorderTileBackground "white"
    TitleForeground "black"
    TitleBackground "white"
}

# I created beifang.bmp with the program "bitmap". Here I tell twm to
# use it as the default highlight pattern on windows' title bars:
Pixmap
{
    TitleHighlight "/home/larry/x/beifang.bmp"
}

# Don't worry about this stuff, it's only for power users :-))
BorderWidth 2
TitleFont "-adobe-new century schoolbook-bold-r-normal--14-140-75-75-p-87-iso8859-1"
MenuFont "6x13"
IconFont "lucidasans-italic-14"
ResizeFont "fixed"
Zoom 50
RandomPlacement

# These programs will not get a window titlebar by default:
NoTitle
{
    "stamp"
    "xload"
    "xclock"
    "xlogo"
    "xbiff"
    "xeyes"
    "oclock"
    "xoid"
}

# "AutoRaise" means that a window is brought to the front whenever the
# mouse pointer enters it. I find this annoying, so I have it turned

```



```

# off. As you can see, I inherited my .twmrc from people who also did
# not like autoraise.
AutoRaise
{
  "nothing"      # I don't like auto-raise # Me either # nor I
}

# Here is where the mouse button functions are defined. Notice the
# pattern: a mouse button pressed on the root window, with no modifier
# key being pressed, always brings up a menu. Other locations usually
# result in window manipulation of some kind, and modifier keys are
# used in conjunction with the mouse buttons to get at the more
# sophisticated window manipulations.
#
# You don't have to follow this pattern in your own .twmrc -- it's
# entirely up to you how you arrange your environment.

# Button = KEYS : CONTEXT : FUNCTION
# -----
Button1 =      : root      : f.menu "main"
Button1 =      : title     : f.raise
Button1 =      : frame     : f.raise
Button1 =      : icon      : f.iconify
Button1 = m    : window    : f.iconify

Button2 =      : root      : f.menu "stuff"
Button2 =      : icon      : f.move
Button2 = m    : window    : f.move
Button2 =      : title     : f.move
Button2 =      : frame     : f.move
Button2 = s    : frame     : f.zoom
Button2 = s    : window    : f.zoom

Button3 =      : root      : f.menu "x"
Button3 =      : title     : f.lower
Button3 =      : frame     : f.lower
Button3 =      : icon      : f.raiselower

# You can write your own functions; this one gets used in the menu
# "windowops" near the end of this file:
Function "raise-n-focus"
{
  f.raise
  f.focus
}

# Okay, below are the actual menus referred to in the mouse button
# section). Note that many of these menu entries themselves call

```

```

# sub-menus. You can have as many levels of menus as you want, but be
# aware that recursive menus don't work. I've tried it.

menu "main"
{
  "Vanilla"      f.title
  "Emacs"        f.menu "emacs"
  "Logins"       f.menu "logins"
  "Ilock"        f.menu "ilock"
  "Misc"         f.menu "misc"
}

# This allows me to invoke emacs on several different machines. See
# the section on .rhosts files for more information about how this
# works:
menu "emacs"
{
  "Emacs"        f.title
  "here"         !"/usr/bin/emacs &"
  ""             f.nop
  "phylo"        !"rsh phylo \"emacs -d floss:0\" &"
  "geta"         !"rsh geta \"emacs -d floss:0\" &"
  "darwin"       !"rsh darwin \"emacs -d floss:0\" &"
  "ninja"        !"rsh ninja \"emacs -d floss:0\" &"
  "indy"         !"rsh indy \"emacs -d floss:0\" &"
  "oberlin"     !"rsh cs.oberlin.edu \"emacs -d floss.life.uiuc.edu:0\" &"
  "gau"         !"rsh gate-1.gau.ai.mit.edu \"emacs -d floss.life.uiuc.edu:0\" &"
}

# This allows me to invoke xterms on several different machines. See
# the section on .rhosts files for more information about how this
# works:
menu "logins"
{
  "Logins"       f.title
  "here"         !"/usr/bin/111/xterm -ls -T 'hostname' -n 'hostname' &"
  "phylo"        !"rsh phylo \"xterm -ls -display floss:0 -T phylo\" &"
  "geta"         !"rsh geta \"xterm -ls -display floss:0 -T geta\" &"
  "darwin"       !"rsh darwin \"xterm -ls -display floss:0 -T darwin\" &"
  "ninja"        !"rsh ninja \"xterm -ls -display floss:0 -T ninja\" &"
  "indy"         !"rsh indy \"xterm -ls -display floss:0 -T indy\" &"
}

# The ilock screensaver, called with various options (each of which
# gives a different pretty picture):
menu "ilock"
{
  "Hop"         !"xlock -mode hop &"
}

```

```

"Qix" !"xlock -mode qix &"
"Flame" !"xlock -mode flame &"
"Worm" !"xlock -mode worm &"
"Swarm" !"xlock -mode swarm &"
"Hop NL" !"xlock -mode hop -nolock &"
"Qix NL" !"xlock -mode qix -nolock &"
"Flame NL" !"xlock -mode flame -nolock &"
"Worm NL" !"xlock -mode worm -nolock &"
"Swarm NL" !"xlock -mode swarm -nolock &"
}

# Miscellaneous programs I run occasionally:
menu "misc"
{
"Xload" !"/usr/bin/X11/xload &"
"XV" !"/usr/bin/X11/xv &"
"Bitmap" !"/usr/bin/X11/bitmap &"
"Tetris" !"/usr/bin/X11/xtetris &"
"Hextris" !"/usr/bin/X11/xhextris &"
"XRoach" !"/usr/bin/X11/xroach &"
"Analog Clock" !"/usr/bin/X11/xclock -analog &"
"Digital Clock" !"/usr/bin/X11/xclock -digital &"
}

# This is the one I bound to the middle mouse button:
menu "stuff"
{
"Chores" f.title
"Sync" !"/bin/sync"
"Who" !"who | xmessage -file - -columns 80 -lines 24 &"
"Xhost +" !"/usr/bin/X11/xhost + &"
"Rootclear" !"/home/larry/bin/rootclear &"
}

# X functions that are sometimes convenient:
menu "x"
{
"X Stuff" f.title
"Xhost +" !"xhost + &"
"Refresh" f.refresh
"Source .twmrc" f.twmrc
"(De)Iconify" f.iconify
"Move Window" f.move
"Resize Window" f.resize
"Destroy Window" f.destroy
"Window Ops" f.menu "windovops"
"" f.nop
"Kill twm" f.quit
}

```

```

}

# This is a submenu from above:
menu "windovops"
{
"Window Ops" f.title
"Show Icon Mgr" f.showiconmgr
"Hide Icon Mgr" f.hideiconmgr
"Refresh" f.refresh
"Refresh Window" f.winrefresh
"twm version" f.version
"Focus on Root" f.unfocus
"Source .twmrc" f.twmrc
"Cut File" f.cutfile
"(De)Iconify" f.iconify
"DeIconify" f.deiconify
"Move Window" f.move
"ForceMove Window" f.forcemove
"Resize Window" f.resize
"Raise Window" f.raise
"Lower Window" f.lower
"Raise or Lower" f.raiselower
"Focus on Window" f.focus
"Raise-n-Focus" f.function "raise-n-focus"
"Destroy Window" f.destroy
"Kill twm" f.quit
}

```

Whew! Believe me, that's not even the most involved .twmrc I've ever seen. It's quite probable that some decent example .twmrc files came with your X. Take a look in the directory /usr/lib/X11/twm/ or /usr/X11/lib/X11/twm and see what's there.

One bug to watch out for with .twmrc files is forgetting to put the & after a command on a menu. If you notice that X just freezes when you run certain commands, chances are that this is the cause. Break out of X with **Control**-**Alt**-**Backspace**, edit your .twmrc, and try again.

9.2.2 Fvwm Configuration

If you are using fvwm, the directory /usr/lib/X11/fvwm/ (or /usr/X11/lib/X11/fvwm/) has some good example config files in it, as well.

[Folks: I don't know anything about fvwm, although I might be able to grok something from the example config files. Then again, so could the reader :-). Also, given the decent but small system.twmrc in the above-mentioned directory, I wonder if it's worth it for me to provide that lengthy example with my own .twmrc. It's in for now, but I don't know whether we want to leave it there or not. -Karl]

9.3 Other Init Files

Some other initialization files of note are:

```
.emacs    Read by the Emacs text editor when it starts up.
.netrc    Gives default login names and passwords for ftp.
.rhosts   Makes your account remotely accessible.
.forward  For automatic mail forwarding.
```

9.3.1 The Emacs Init File

If you use `emacs` as your primary editor, then the `.emacs` file is quite important. It is dealt with at length in Chapter 8.

9.3.2 FTP Defaults

Your `.netrc` file allows you to have certain `ftp` defaults set before you run `ftp`. Here is a small sample `.netrc`:

```
machine floss.life.uiuc.edu login larry password fishSticks
machine darwin.life.uiuc.edu login larry password fishSticks
machine geta.life.uiuc.edu login larry password fishSticks
machine phylo.life.uiuc.edu login larry password fishSticks
machine ninja.life.uiuc.edu login larry password fishSticks
machine indy.life.uiuc.edu login larry password fishSticks

machine clone.mcs.anl.gov login fogel password doorm0
machine osprey.mcs.anl.gov login fogel password doorm0
machine tern.mcs.anl.gov login fogel password doorm0
machine altair.mcs.anl.gov login fogel password doorm0
machine dalek.mcs.anl.gov login fogel password doorm0
machine juju.mcs.anl.gov login fogel password doorm0

machine sunsite.unc.edu login anonymous password larry@cs.oberlin.edu
```

Each line of your `.netrc` specifies a machine name, a login name to use by default for that machine, and a password. This is a great convenience if you do a lot of `ftp`-ing and are tired of constantly typing in your username and password at various sites. The `ftp` program will try to log you in automatically using the information found in your `.netrc` file, if you `ftp` to one of the machines listed in the file.

You can tell `ftp` to ignore your `.netrc` and not attempt auto-login by invoking it with the `-n` option: `"ftp -n"`.

You must make sure that your `.netrc` file is readable *only* by you. Use the `chmod` program to set the file's read permissions. If other people can read it, that means they can find out your password

at various other sites. This is about as big a security hole as one can have; to encourage you to be careful, `ftp` and other programs that look for the `.netrc` file will actually refuse to work if the read permissions on the file are bad.

There's more to the `.netrc` file than what I've said; when you get a chance, do `"man .netrc"` or `"man ftp"`.

9.3.3 Allowing Easy Remote Access to Your Account

If you have an `.rhosts` file in your home directory, it will allow you to run programs on this machine remotely. That is, you might be logged in on the machine `cs.oberlin.edu`, but with a correctly configured `.rhosts` file on `floss.life.uiuc.edu`, you could run a program on `floss.life.uiuc.edu` and have the output go to `cs.oberlin.edu`, without ever having to log in or type a password.

A `.rhosts` file looks like this:

```
frobnozz.cs.knowledge.edu jsmith
aphrodite.classics.hahvaahd.edu uphilps
frobbo.hoola.com trixie
```

The format is fairly straightforward: a machine name, followed by username. Suppose that that example is in fact my `.rhosts` file on `floss.life.uiuc.edu`. That would mean that I could run programs on `floss`, with output going to any of the machines listed, as long as I were also logged in as the corresponding user given for that machine when I tried to do it.

The exact mechanism by which one runs a remote program is usually the `rsh` program. It stands for "remote shell", and what it does is start up a shell on a remote machine and execute a specified command. For example:

```
frobbo$ whoami
trixie
frobbo$ rsh floss.life.uiuc.edu "ls ~"
foo.txt mbox url.ps snax.txt
frobbo$ rsh floss.life.uiuc.edu "more ~/snax.txt"
[snax.txt comes paging by here]
```

User `trixie` at `floss.life.uiuc.edu`, who had the example `.rhosts` shown previously, explicitly allows `trixie` at `frobbo.hoola.com` to run programs as `trixie` from `floss`.

You don't have to have the same username on all machines to make a `.rhosts` work right. Use the `-l` option to `rsh`, to tell the remote machine what username you'd like to use for logging in. If that username exists on the remote machine, and has a `.rhosts` file with your current (i.e.: local) machine and username in it, then your `rsh` will succeed.

```
frobbo$ whoami
trixie
frobbo$ rsh -l larry floss.life.uiuc.edu "ls ~"
[Insert a listing of my directory on floss here]
```

This will work if user larry on floss.life.uiuc.edu has a .rhosts file which allows trixie from frobbo.hoopla.com to run programs in his account. Whether or not they are the same person is irrelevant: the only important things are the usernames, the machine names, and the entry in larry's .rhosts file on floss. Note that trixie's .rhosts file on frobbo doesn't enter into it, only the one on the remote machine matters.

There are other combinations that can go in a .rhosts file—for example, you can leave off the username following a remote machine name, to allow any user from that machine to run programs as you on the local machine! This is, of course, a security risk: someone could remotely run a program that removes your files, just by virtue of having an account on a certain machine. If you're going to do things like leave off the username, then you ought to make sure that your .rhosts file is readable by you and no one else.

9.3.4 Mail Forwarding

You can also have a .forward file, which is not strictly speaking an "init file". If it contains an email address, then all mail to you will be forwarded to that address instead. This is useful when you have accounts on many different systems, but only want to read mail at one location.

There is a host of other possible initialization files. The exact number will vary from system to system, and is dependent on the software installed on that system. One way to learn more is to look at files in your home directory whose names begin with ".". These files are not all guaranteed to be init files, but it's a good bet that most of them are.

9.4 Seeing Some Examples

The ultimate example I can give you is a running Linux system. So, if you have Internet access, feel free to telnet to floss.life.uiuc.edu. Log in as "guest", password "explorer", and poke around. Most of the example files given here can be found in /home/kfogel, but there are other user directories as well. You are free to copy anything that you can read. Please be careful: floss is not a terribly secure box, and you can almost certainly gain root access if you try hard enough. I prefer to rely on trust, rather than constant vigilance, to maintain security.

Chapter 11

Funny Commands

Well, most people who had to do with the UNIX commands exposed in this chapter will not agree with this title. "What the heck! You have just shown me that the Linux interface is very standard, and now we have a bunch of commands, each one working in a completely different way. I will never remember all those options, and you are saying that they are *funny*?" Yes, you have just seen an example of hackers' humor. Besides, look at it from the bright side: there is no MS-DOS equivalent of these commands. If you need them, you have to purchase them, and you never know how their interface will be. Here they are a useful – and inexpensive – add-on, so enjoy!

The set of commands dwelled on in this chapter covers `find`, which lets the user search in the directory tree for specified groups of files; `tar`, useful to create some archive to be shipped or just saved; `dd`, the low-level copier; and `sort`, which ... yes, sorts files. A last proviso: these commands are by no means standardized, and while a core of common options could be found on all *IX systems, the (GNU) version which is explained below, and which you can find in your Linux system, has usually many more capabilities. So if you plan to use other UNIX-like operating systems, please don't forget to check their man page in the target system to learn the maybe not-so-little differences.

11.1 `find`, the file searcher

11.1.1 Generalities

Among the various commands seen so far, there were some which let the user recursively go down the directory tree in order to perform some action: the canonical examples are `ls -R` and `rm -R`. Good. `find` is *the* recursive command. Whenever you are thinking "Well, I have to do so-and-so on all those kind of files in my own partition", you have better think about using `find`. In a certain sense the fact that `find` finds files is just a side effect: its real occupation is to evaluate.

The basic structure of the command is as follows:

```
find path [...] expression [...]
```

This at least on the GNU version; other version do not allow to specify more than one path, and besides it is very uncommon the need to do such a thing. The rough explanation of the command syntax is rather simple: you say from where you want to start the search (the *path* part; with GNU `find` you can omit this and it will be taken as default the current directory `.`), and which kind of search you want to perform (the *expression* part).

The standard behavior of the command is a little tricky, so it's worth to note it. Let's suppose that in your home directory there is a directory called `garbage`, containing a file `foobar`. You happily type `find . -name foobar` (which as you can guess searches for files named `foobar`), and you obtain ... nothing else than the prompt again. The trouble lies in the fact that `find` is by default a silent command; it just returns 0 if the search was completed (with or without finding anything) or a non-zero value if there had been some problem. This does not happen with the version you can find on Linux, but it is useful to remember it anyway.

11.1.2 Expressions

The *expression* part can be divided itself in four different groups of keywords: *options*, *tests*, *actions*, and *operators*. Each of them can return a true/false value, together with a side effect. The difference among the groups is shown below.

options affect the overall operation of `find`, rather than the processing of a single file. An example is `-follow`, which instructs `find` to follow symbolic links instead of just stating the inode. They always return true.

tests are real tests (for example, `-empty` checks whether the file is empty), and can return true or false.

actions have also a side effect the name of the considered file. They can return true or false too.

operators do not really return a value (they can conventionally be considered as true), and are used to build complex expression. An example is `-or`, which takes the logical OR of the two subexpressions on its side. Notice that when juxtaposing expression, a `-and` is implied.

Note that `find` relies upon the shell to have the command line parsed; it means that all keyword must be embedded in white space and especially that a lot of nice characters have to be escaped, otherwise they would be mangled by the shell itself. Each escaping way (backslash, single and double quotes) is OK; in the examples the single character keywords will be usually quoted with backslash, because it is the simplest way (at least in my opinion. But it's me who is writing these notes!)

11.1.3 Options

Here there is the list of all options known by GNU version of `find`. Remember that they always return true.

- `-daystart` measures elapsed time not from 24 hours ago but from last midnight. A true hacker probably won't understand the utility of such an option, but a worker who programs from eight to five does appreciate it.
- `-depth` processes each directory's contents before the directory itself. To say the truth, I don't know many uses of this, apart for an emulation of `rm -F` command (of course you cannot delete a directory before all files in it are deleted too ...)
- `-follow` defences (that is, follows) symbolic links. It implies option `-noleaf`; see below.
- `-noleaf` turns off an optimization which says "A directory contains two fewer subdirectories than their hard link count". If the world were perfect, all directories would be referenced by each of their subdirectories (because of the `..` option), as `.` inside itself, and by its "real" name from its parent directory.
That means that every directory must be referenced at least twice (once by itself, once by its parent) and any additional references are by subdirectories. In practice however, symbolic links and distributed filesystems¹ can disrupt this. This option makes `find` run slightly slower, but may give expected results.
- `-maxdepth levels`, `-mindepth levels`, where `levels` is a non-negative integer, respectively say that at most or at least `levels` levels of directories should be searched. A couple of examples is mandatory: `-maxdepth 0` indicates that the command should be performed just on the arguments in the command line, i.e., without recursively going down the directory tree; `-mindepth 1` inhibits the processing of the command for the arguments in the command line, while all other files down are considered.
- `-version` just prints the current version of the program.
- `-xdev`, which is a misleading name, instructs `find` not to cross device, i.e. changing filesystem. It is very useful when you have to search for something in the root filesystem; in many machines it is a rather small partition, but a `find /` would otherwise search the whole structure!

11.1.4 Tests

The first two tests are very simple to understand: `-false` always return false, while `-true` always return true. Other tests which do not need the specification of a value are `-empty`, which returns true whether the file is empty, and the couple `-nouser / -nogroup`, which return true in the case that no entry in `/etc/passwd` or `/etc/group` match the user/group id of the file owner. This is a common thing which happens in a multiuser system; a user is deleted, but files owned by her remain in the strangest part of the filesystems, and due to Murphy's laws take a lot of space.

Of course, it is possible to search for a specific user or group. The tests are `-uid nn` and `-gid nn`. Unfortunately it is not possible to give directly the user name, but it is necessary to use the numeric id, `nn`.

¹Distributed filesystems allow files to appear like their local to a machine when they are actually located somewhere else.

allowed to use the forms `+nn`, which means "a value strictly greater than `nn`", and `-nn`, which means "a value strictly less than `nn`". This is rather silly in the case of UIDs, but it will turn handy with other tests.

Another useful option is `-type c`, which returns true if the file is of type `c`. The mnemonics for the possible choices are the same found in `ls`; so we have `b` when the file is a block special; `c` when the file is character special; `d` for directories; `p` for named pipes; `l` for symbolic links, and `s` for sockets. Regular files are indicated with `f`. A related test is `-xtype`, which is similar to `-type` except in the case of symbolic links. If `-follow` has not been given, the file pointed at is checked, instead of the link itself. Completely unrelated is the test `-fstype type`. In this case, the filesystem type is checked. I think that the information is got from file `/etc/mstab`, the one stating the mounting filesystems; I am certain that types `nfs`, `tmp`, `msdos` and `ext2` are recognized.

Tests `-inum nn` and `-links nn` check whether the file has inode number `nn`, or `nn` links, while `-size nn` is true if the file has `nn` 512-bytes blocks allocated. (well, not precisely: for sparse files unallocated blocks are counted too). As nowadays the result of `ls -s` is not always measured in 512-bytes chunks (Linux for example uses 1k as the unit), it is possible to append to `nn` the character `b`, which means to count in bytes, or `k`, to count in kilobytes.

Permission bits are checked through the test `-perm mode`. If `mode` has no leading sign, then the permission bits of the file must exactly match them. A leading `-` means that all permission bits must be set, but makes no assumption for the other; a leading `+` is satisfied just if any of the bits are set. Oops! I forgot saying that the mode is written in octal or symbolically, like you use them in `chmod`.

Next group of tests is related to the time in which a file has been last used. This comes handy when a user has filled his space, as usually there are many files he did not use since ages, and whose meaning he has forgot. The trouble is to locate them, and `find` is the only hope in sight. `-atime nn` is true if the file was last accessed `nn` days ago, `-ctime nn` if the file status was last changed `nn` days ago for example, with a `chmod` and `-mtime nn` if the file was last modified `nn` days ago. Sometimes you need a more precise timestamp; the test `-newer file` is satisfied if the file considered has been modified later than `file`. So, you just have to use `touch` with the desired date, and you're done. GNU `find` add the tests `-anewer` and `-cnewer` which behave similarly; and the tests `-amin`, `-cmin` and `-mmin` which count time in minutes instead than 24-hours periods.

Last but not the least, the test I use more often. `-name pattern` is true if the file name exactly matches `pattern`, which is more or less the one you would use in a standard `ls`. Why 'more or less'? Because of course you have to remember that all the parameters are processed by the shell, and those lovely metacharacters are expanded. So, a test like `-name foo*` won't return what you want, and you should either write `-name foo` or `-name "foo*"`. This is probably one of the most common mistakes made by careless users, so write it in BIG letters on your screen. Another problem is that, like with `ls`, leading dots are not recognized. To cope with this, you can use test `-path pattern` which does not worry about dot and slashes when comparing the path of the considered file with `pattern`.

11.1.5 Actions

I have said that actions are those which actually do something. Well, `-prune` rather does not do something, i.e. descending the directory tree (unless `-depth` is given). It is usually `find` together with `-fstype`, to choose among the various filesystems which should be checked.

The other actions can be divided into two broad categories;

- Actions which *print* something. The most obvious of these – and indeed, the default action of `find` – is `-print` which just print the name of the file(s) matching the other conditions in the command line, and returns true. A simple variants of `-print` is `-print file`, which uses `file` instead of standard output, `-ls` lists the current file in the same format as `ls -dils`; `-printf format` behaves more or less like C function `printf()`, so that you can specify how the output should be formatted, and `-printf file format` does the same, but writing on `file`. These action too return true.
- Actions which *execute* something. Their syntax is a little odd and they are used widely, so please look at them.

`-exec command \`; the command is executed, and the action returns true if its final status is 0, that is regular execution of it. The reason for the `\`; is rather logical: `find` does not know where the command ends, and the trick to put the `exec` action at the end of the command is not applicable. Well, the best way to signal the end of the command is to use the character used to do this by the shell itself, that is `'`; but of course a semicolon all alone on the command line would be eaten by the shell and never sent to `find`, so it has to be escaped. The second thing to remember is how to specify the name of the current file within `command`, as probably you did all the trouble to build the expression to do something, and not just to print data. This is done by means of the string `{}`. Some old versions of `find` require that it must be embedded in white space – not very handy if you needed for example the whole path and not just the file name – but with GNU `find` could be anywhere in the string composing `command`. And shouldn't it be escaped or quoted, you surely are asking? Amazingly, I never had to do this neither under `tcsh` nor under `bash` (`sh` does not consider `{}` and `}` as special characters, so it is not much of a problem). My idea is that the shells “know” that `{}` is not an option making sense, so they do not try to expand them, luckily for `find` which can obtain it untouched.

`-ok command \`; behaves like `-exec`, with the difference that for each selected file the user is asked to confirm the command; if the answer starts with `y` or `Y`, it is executed, otherwise not, and the action returns false.

11.1.6 Operators

There are a number of operators; here there is a list, in order of decreasing precedence.

`(expr)`

forces the precedence order. The parentheses must of course be quoted, as they are meaningful for the shell too.

`! expr`

`-not expr`

change the truth value of expression, that is if `expr` is true, it becomes false. The exclamation mark needn't be escaped, because it is followed by a white space.

`expr1 expr2`

`expr1 -a expr2`

`expr1 -and expr2`

all correspond to the logical AND operation, which in the first and most common case is implied. `expr2` is not evaluated, if `expr1` is false.

`expr1 -o expr2`

`expr1 -or expr2`

correspond to the logical OR operation. `expr2` is not evaluated, if `expr1` is true.

`expr1 , expr2`

is the list statement; both `expr1` and `expr2` are evaluated (together with all side effects, of course!), and the final value of the expression is that of `expr2`.

11.1.7 Examples

Yes, `find` has just too many options, I know. But there are a lot of cooked instances which are worth to remember, because they are used very often. Let's see some of them.

```
% find . -name foo\* -print
```

finds all file names starting with `foo`. If the string is embedded in the name, probably it is more sensitive to write something like `"*foo*"`, rather than `foo`.

```
% find /usr/include -xtype f -exec grep foobar \
/dev/null {} \;
```

is a `grep` executed recursively starting from directory `/usr/include`. In this case, we are interested both in regular file and in symbolic links which point to regular files, hence the `-xtype` test. Many times it is simpler to avoid specifying it, especially if we are rather sure no binary file contains the wanted string. And why the `/dev/null` in the command? It's a trick to force `grep` to write the file name where a match has been found. The command `grep` is applied to each file in a different invocation, and so it doesn't think it is necessary to output the file name. But now there are two files, i.e. the current one and `/dev/null`! Another possibility should be to pipe the command to `xargs` and let it perform the `grep`. I just tried it, and completely smashed my filesystem (together with these notes which I am trying to recover by hand :-)).

```
% find / -atime +1 -fstype ext2 -name core \
-exec rm {} \;
```

is a classical job for `crontab`. It deletes all file named `core` in filesystems of type `ext2` which have not been accessed in the last 24 hours. It is possible that someone wants to use the `core` file to perform a post mortem dump, but nobody could remember what he was doing after 24 hours...

```
% find /home -xdev -size +500k -ls > piggies
```

is useful to see who has those files who clog the filesystem. Note the use of `-xdev`; as we are interested in just one filesystem, it is not necessary to descend other filesystems mounted under `/home`.

11.1.8 A last word

Keep in mind that `find` is a very time consuming command, as it has to access each and every inode of the system in order to perform its operation. It is therefore wise to combine how many operations you need in a unique invocation of `find`, especially in the 'housekeeping' jobs usually ran via a `crontab` job. A enlightening example is the following: let's suppose that we want to delete files ending in `.BAK` and change the protection of all directories to `771` and that of all files ending in `.sh` to `755`. And maybe we are mounting NFS filesystems on a dial-up link, and we'd like not to check for files there. Why writing three different commands? The most effective way to accomplish the task is this:

```
% find . \( -fstype nfs -prune \) -o \
  \( -type d      -a -exec chmod 771 {} \; \) -o \
  \( -name "*.BAK" -a -exec /bin/rm {} \; \) -o \
  \( -name "*.sh"  -a -exec chmod 755 {} \; \)
```

It seems ugly (and with much abuse of backslashes!), but looking closely at it reveals that the underlying logic is rather straightforward. Remember that what is really performed is a true/false evaluation; the embedded command is just a side effect. But this means that it is performed only if `find` must evaluate the `exec` part of the expression, that is only if the left side of the subexpression evaluates to true. So, if for example the file considered at the moment is a directory then the first `exec` is evaluated and the permission of the inode is changed to `771`; otherwise it forgets all and steps to the next subexpression. Probably it's easier to see it in practice than to writing it down; but after a while, it will become a natural thing.

11.2 tar, the tape archiver

11.2.1 Introduction

11.2.2 Main options

11.2.3 Modifiers

11.2.4 Examples

11.3 dd, the data duplicator

Legend says that back in the mists of time, when the first UNIX was created, its developers needed a low level command to copy data between devices. As they were in a hurry, they decided to borrow the syntax used by IBM-360 machines, and to develop later an interface consistent with that of the other commands. Time passed, and all were so used with the odd way of using `dd` that it stuck. I don't know whether it is true, but it is a nice story to tell.

11.3.1 Options

To say the truth, `dd` it's not completely unlike the other Unix command: it is indeed a filter, that is it reads by default from the standard input and writes to the standard output. So if you just type `dd` at the terminal it remains quiet, waiting for input, and a `ctrl-C` is the only sensitive thing to type.

The syntax of the command is as follows:

```
dd [if=file] [of=file] [ibs=bytes] [obs=bytes] [bs=bytes] [cbs=bytes] [skip=blocks]
   [seek=blocks] [count=blocks] [conv={ascii, ebedic, ibm, block, unblock,
   lease, ucase, swab, noerror, notrunc, sync}]
```

All options are of the form *option=value*. No space is allowed either before or after the equal sign; this used to be annoying, because the shell did not expand a filename in this situation, but the version of `bash` present in Linux is rather smart, so you don't have to worry about that. It is important also to remember that all numbered values (`bytes` and `blocks` above) can be followed by a multiplier. The possible choices are `b` for block, which multiplies by 512, `k` for kilobytes (1024), `w` for word (2), and `m` multiplies by `m`.

The meaning of options is explained below.

- `if=filein` and `of=fileout` instruct `dd` to respectively read from *filein* and write to *fileout*. In the latter case, the output file is truncated to the value given to `seek`, or if the keyword is not

present, to 0 (that is deleted), before performing the operation. But look below at option `notrunc`.

- `ibs=nn` and `obs=nn` specify how much bytes should be read or write at a time. I think that the default is 1 block, i.e. 512 bytes, but I am not very sure about it: certainly it works that way with plain files. These parameters are very important when using special devices as input or output; for example, reading from the net should set `ibs` at 10k, while a high density 3.5" floppy has as its natural block size 18k. Failing to set these values could result not only in longer time to perform the command, but even in timeout errors, so be careful.
- `bs=nn` both reads and writes `nn` bytes at a time. It overrides `ibs` and `obs` keywords.
- `cbs=nn` sets the conversion buffers to `nn` bytes. This buffer is used when translating from ASCII to EBCDIC, or from an unblocked device to a blocked one. For example, files created under VMS have often a block size of 512, so you have to set `cbs` to 1b when reading a foreign VMS tape. Hope that you don't have to mess with these things!
- `skip=nb1` and `seek=nb1` tell the program to skip `nb1` blocks respectively at the beginning of input and at the beginning of output. Of course the latter case makes sense if conversion `notrunc` is given, see below. Each block's size is the value of `ibs` (`obs`). Beware: if you did not set `ibs` and write `skip=1b` you are actually skipping 512x512 bytes, that is 256KB. It was not precisely what you wanted, wasn't it?
- `count=nb1` means to copy only `nb1` blocks from input, each of the size given by `ibs`. This option, together with the previous, turns useful if for example you have a corrupted file and you want to recover how much it is possible from it. You just skip the unreadable part and get what remains.
- `conv=conversion,[conversion...]` convert the file as specified by its argument. Possible conversions are `ascii`, which converts from EBCDIC to ASCII; `ebcdic` and `ibm`, which both perform an inverse conversion (yes, there is not a unique conversion from EBCDIC to ASCII! The first is the standard one, but the second works better when printing files on a IBM printer); `block`, which pads newline-terminated records to the size of `cbs`, replacing newline with trailing spaces; `unblock`, which performs the opposite (eliminates trailing spaces, and replaces them with newline); `lcase` and `ucase`, to convert text to lowercase and uppercase; `swab`, which swaps every pair of input bytes (for example, to use a file containing short integers written on a 680x0 machine in an Intel-based machine you need such a conversion); `noerror`, to continue processing after read errors; `sync`, which pads input block to the size of `ibs` with trailing NULs.

11.3.2 Examples

The canonical example is the one you have probably bumped at when you tried to create the first Linux diskette: how to write to a floppy without a MS-DOS filesystem. The solution is simple:

```
% dd if=disk.img of=/dev/fd0 obs=18k count=80
```

I decided not to use `ibs` because I don't know which is the better block size for a hard disk, but in this case no harm would have been if instead of `obs` I use `bs` - it could even be a trifle quicker. Notice the explicitation of the number of sectors to write (18KB is the occupation of a sector, so `count` is set to 80) and the use of the low-level name of the floppy device.

Another useful application of `dd` is related to the network backup. Let's suppose that we are on machine *alpha* and that on machine *beta* there is the tape unit `/dev/rst0` with a tar file we are interested in getting. We have the same rights on both machines, but there is no space on *beta* to dump the tar file. In this case, we could write

```
% rsh beta 'dd if=/dev/rst0 ibs=8k obs=20k' | tar xvBf -
```

to do in a single pass the whole operation. In this case, we have used the facilities of `rsh` to perform the reading from the tape. Input and output sizes are set to the default for these operations, that is 8KB for reading from a tape and 20KB for writing to ethernet; from the point of view of the other side of the tar, there is the same flow of bytes which could be got from the tape, except the fact that it arrives in a rather erratic way, and the option `B` is necessary.

I forgot: I don't think at all that `dd` is an acronym for "data duplicator", but at least this is a nice way to remember its meaning ...

11.4 sort, the data sorter

11.4.1 Introduction

11.4.2 Options

11.4.3 Examples