A/UX Programming Languages and Tools

Volume 2

Release 3.0

Limited Warranty on Media and Replacement

If you discover physical defects in the manuals distributed with an Apple product or in the media on which a software product is distributed, Apple will replace the media or manuals at no charge to you, provided you return the item to be replaced with proof of purchase to Apple or an authorized Apple dealer during the 90-day period after you purchased the software. In addition, Apple will replace damaged software media and manuals for as long as the software product is included in Apple's Media Exchange Program. While not an upgrade or update method, this program offers additional protection for up to two years or more from the date of your original purchase. See your authorized Apple dealer for program coverage and details. In some countries the replacement period may be different; check with your authorized Apple dealer.

ALL IMPLIED WARRANTIES ON THE MEDIA AND MANUALS, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, ARE LIMITED IN DURATION TO NINETY (90) DAYS FROM THE DATE OF THE ORIGINAL RETAIL PURCHASE OF THIS PRODUCT.

Even though Apple has tested the software and reviewed the documentation, **APPLE MAKES NO WARRANTY OR REPRESENTATION, EITHER EXPRESS, OR IMPLIED, WITH RESPECT TO SOFTWARE, ITS QUALITY, PERFORMANCE, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. AS A RESULT, THIS SOFTWARE IS SOLD "AS IS," AND YOU, THE PURCHASER, ARE ASSUMING THE ENTIRE RISK AS TO ITS QUALITY AND PERFORMANCE.**

IN NO EVENT WILL APPLE BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES RESULTING FROM ANY DEFECT IN THE SOFTWARE OR ITS DOCUMENTATION, even if advised of the possibility of such damages. In particular, Apple shall have no liability for any programs or data stored in or used with Apple products, including the costs of recovering such programs or data.

THE WARRANTY AND REMEDIES SET FORTH ABOVE ARE EXCLUSIVE AND IN LIEU OF ALL OTHERS, ORAL OR WRITTEN, EXPRESS, OR IMPLIED. No Apple dealer, agent, or employee is authorized to make any modification, extension, or addition to this warranty.

Some states do not allow the exclusion or limitation of implied warranties or liability for incidental or consequential damages, so the above limitation or exclusion may not apply to you. This warranty gives you specific legal rights, and you may also have other rights which vary from state to state.

Apple Computer, Inc.

© 1992, Apple Computer, Inc. and UniSoft Corporation. All rights reserved.

Portions of this document have been previously copyrighted by AT&T Information Systems and the Regents of the University of California, and are reproduced with permission. Under the copyright laws, this manual may not be copied, in whole or part, without the written consent of Apple or UniSoft. The same proprietary and copyright notices must be affixed to any permitted copies as were affixed to the original. Under the law, copying includes translating into another language or format.

The Apple logo is a registered trademark of Apple Computer, Inc. Use of the "keyboard" logo (Option-Shift-K) for commercial purposes without the prior written consent of Apple may constitute trademark infringement and unfair competition in violation of federal and state laws.

Apple Computer, Inc.
20525 Mariani Avenue
Cupertino, CA 95014-6299
(408) 996-1010

Apple, the Apple logo, APDA, AppleShare, AppleTalk, A/UX, EtherTalk, ImageWriter, LaserWriter, LocalTalk, Macintosh, MacTCP, MPW, MultiFinder, SANE, and TokenTalk are trademarks of Apple Computer, Inc., registered in the United States and other countries.

Apple Desktop Bus, Finder, MacX, QuickDraw, ResEdit, and SuperDrive are trademarks of Apple Computer, Inc.

Adobe, Adobe Illustrator, and PostScript are registered trademarks of Adobe Systems Incorporated.

cdb is a trademark of Third Eye Software, Inc.

DEC, Internet, PDP-11, VAX, and VT100 are trademarks of Digital Equipment Corporation.

Electrocomp 2000 is a trademark of Image Graphics, Inc.

IBM is a registered trademark, and System 370 is a trademark, of International Business Machines Corporation.

ITC Garamond and ITC Zapf Dingbats are registered trademarks of International Typeface Corporation.

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

Motorola is a registered trademark of Motorola Corporation.

NFS, SPARC, and SUN are trademarks of Sun Microsystems, Inc.

NuBus is a trademark of Texas Instruments.

QuarkXPress is a registered trademark of Quark, Inc.

UNIX is a registered trademark of UNIX System Laboratories, Inc.

X Window System is a trademark of Massachusetts Institute of Technology.

Zilog is a registered trademark of Zilog, Inc.

Simultaneously published in the United States and Canada.

Mention of third-party products is for informational purposes only and constitutes neither an endorsement nor a recommendation. Apple assumes no responsibility with regard to the performance or use of these products.

Contents

Figures, Tables, and Listings / xvii

About This Guide / xxi

Who should use this guide / xxi
What you need to know / xxi
What's covered in this guide / xxii
Where to go for more information / xxii
How to use this guide / xxiii
Conventions used in this guide / xxiii
- Keys and key combinations / xxiii
- Terminology / xxiv
- The `Courier` font / xxiv
- Font styles / xxv
- A/UX command syntax / xxv
- Manual page reference notation / xxvi

For more information / xxvii

1 Overview of Programming Tools / 1-1

Program development tools / 1-2
- Program structure: `cb` / 1-2
- Execution: `cflow` and `prof` / 1-2
- Processing: `m4`, `lex`, and `yacc` / 1-3
- Debugging: `nm` and `od` / 1-3

File manipulation tools / 1-4
File characteristics: `size`, `diff`, and `comm` / 1-4
Maintenance: `make`, SCCS, and `ar` / 1-5
A file-processing language: `awk` / 1-6
Math functions: `dc` and `bc` / 1-6
Screen-oriented tools: `curses` and Commando / 1-6

Part 1 Program Development Tools

2 Programming Tools / 2-1

Improving C program structure: `cb` / 2-2
Generating a C flowgraph: `cflow` / 2-2
Displaying profile data: `prof` / 2-2
A C language preprocessor: `cpp` / 2-3
Finding a function definition quickly: `ctags` / 2-4
Sharing strings from C programs: `xstr` / 2-4
Printing the symbol table for a COFF file: `nm` / 2-5
Obtaining an octal dump of a file: `od` / 2-5

3 `yacc`: A Compiler-Writing System / 3-1

Usage / 3-3
Basic specifications / 3-6
Actions / 3-8
Lexical analysis / 3-12
Parser operation / 3-14
Ambiguity and conflicts / 3-19
Precedence / 3-23
Error handling / 3-27
The `yacc` environment / 3-29
Input style / 3-31
Left recursion / 3-31
Lexical considerations / 3-32
Reserved words / 3-34

Simulating error and accept in actions / 3-34
Accessing values in enclosing rules / 3-34
Arbitrary value types / 3-36
Example: A desk calculator / 3-38
Example: `yacc` input syntax / 3-43
Example: An advanced grammar / 3-46
Backward compatibility / 3-56

4 m4: A Macro Processor / 4-1

Invoking `m4` / 4-3
Defining macros / 4-3
`define` / 4-3
Quoting / 4-5
`changequote` / 4-6
`undefine` / 4-6
`ifdef` / 4-6
Arguments / 4-7
`ifelse` / 4-8
Arithmetic built-ins / 4-9
I/O manipulation / 4-10
`include` and `sinclude` / 4-10
`divert`, `undivert`, and `divnum` / 4-11
`dnl` / 4-14
String manipulation / 4-14
`len` / 4-14
`substr` / 4-15
`index` and `translit` / 4-15
Printing / 4-16
`errprint` / 4-16
`dumpdef` / 4-16
Executing system commands / 4-16
`syscmd` and `maketemp` / 4-16
Interactive use of `m4` / 4-17
Recursive definitions / 4-17
Built-in macro summary / 4-19

5 lex: A Lexical Analyzer / 5-1

Overview of lex usage / 5-3
lex and yacc / 5-4
Program syntax / 5-6
Character set / 5-7
Character classes / 5-7
Arbitrary characters / 5-9
Operators / 5-9
Definitions / 5-10
Repetitions and definitions / 5-12
Rules / 5-12
Regular expressions / 5-12
Optional expressions / 5-13
Repeated expressions / 5-13
Alternation and grouping / 5-14
Context sensitivity / 5-14
Left context sensitivity / 5-15
Flags / 5-16
Start conditions / 5-17
Ambiguous rules / 5-18
Actions / 5-19
The null statement / 5-20
The repetition character / 5-20
printf and ECHO / 5-20
yyleng / 5-21
yymore and yyless / 5-22
lex input and output routines / 5-23
yywrap / 5-24
REJECT / 5-25
Compilation / 5-27
Examples / 5-27
Summary / 5-29

Part 2 File Manipulation Tools

6 File Attribute Tools / 6-1

Comparing source files / 6-2
Finding files: `find` / 6-2
Printing the section sizes of COFF files: `size` / 6-2
Finding the version number of a file: `version` / 6-3
Maintaining portable archives and libraries: `ar` / 6-3

7 `make`: A File Production Tool / 7-1

Using `make` / 7-3
- Writing a makefile / 7-3
- `make` command syntax / 7-5
  - Options / 7-6
  - Using `make` on individual files / 7-8

The description file / 7-8
- Makefile entries / 7-9
  - Targets versus rules / 7-9
  - Built-in targets / 7-10
  - Dependency statements / 7-11
  - Commands / 7-12
- Comments / 7-13
- `include` lines / 7-13
- Macro definitions / 7-13
  - Internal macros / 7-15
  - Dynamic dependency parameters / 7-16
- Options / 7-18
  - Suppressing printing of commands / 7-18
  - Ignoring errors / 7-18
  - Combining commands / 7-19
  - Default commands / 7-19
  - Saving files / 7-20
  - Use of selected options / 7-20

Suffixes and rules / 7-20
- Suffixes / 7-20
- Transformation rules / 7-21
  - The default macro settings / 7-26
- Changing default suffixes and rules / 7-27
  - The default suffix list / 7-27
  - The default rules / 7-28

Operation / 7-28
Environment variables / 7-28
Macros / 7-29
Precedence / 7-35
Macro Testing / 7-37
Attributes / 7-38
Archive libraries / 7-40
SCCS files / 7-42
SCCS filename prefixes / 7-42
SCCS filename suffixes / 7-43
SCCS transformation rules / 7-43
SCCS makefiles / 7-43
Advanced topics / 7-44
Walking the directory tree / 7-44
The `make` predecessor tree / 7-45
The makefile as shell script / 7-46
Unintended targets / 7-46
Mnemonic targets / 7-47
Macro translation / 7-47
Dynamic Include File Dependency Generation / 7-50
A warning for system administrators / 7-52

8 SCCS Reference / 8-1

SCCS for beginners / 8-3
Creating an SCCS file / 8-3
Retrieving a file and storing a new version / 8-4
Retrieving versions / 8-5
On-line information / 8-6
SCCS files / 8-7
Standard A/UX protection / 8-7
SCCS protection mechanisms / 8-8
Administering SCCS / 8-9
Group project administration / 8-9
SCCS file formats / 8-12
SCCS file auditing / 8-12
Delta numbering / 8-13
Branch deltas / 8-14
SCCS command conventions / 8-16
SCCS command arguments / 8-16
Flags / 8-17
Diagnostics / 8-17

Temporary files / 8-17
SCCS ID keywords / 8-20
SCCS command summary / 8-22
Create SCCS files: `admin` / 8-22
SCCS flags / 8-23
Comments and MR numbers / 8-24
Descriptive text / 8-25
Change comments in an SCCS file: `cdc` / 8-26
Combine deltas to save space: `comb` / 8-26
Store a new SCCS file version: `delta` / 8-27
Required temporary files / 8-27
Comments and MR numbers / 8-28
Keywords / 8-29
Removal of temporary files / 8-30
Retrieve an SCCS file version: `get` / 8-30
Retrieving different versions / 8-31
Retrieving a file to create a new delta / 8-32
Concurrent edits of different versions / 8-34
Concurrent edits of the same SID / 8-36
Keyletters that affect output / 8-37
Restore a version unchanged: `unget` / 8-38
On-line explanations: `help` / 8-39
Print parts of an SCCS file: `prs` / 8-39
Remove a specific delta: `rmdel -r` / 8-41
Account for open SCCS files: `sact` / 8-42
Compare two SCCS files: `sccsdiff` / 8-43
Check SCCS file characteristics: `val` / 8-43
Find identifying information: `what` / 8-44

9 `awk` Programming Language / 9-1

`awk` operation / 9-3
Comments / 9-5
Command-line options / 9-6
Invocation modes / 9-7
Interactions with the shell / 9-9
Text input processing / 9-11
Patterns / 9-14
Using expressions for patterns / 9-15
Regular expression syntax / 9-17
`BEGIN` and `END` / 9-19

Actions / 9-20
Components of `awk` programs / 9-21
Flow of control / 9-23
Report generation / 9-27
Reading input: `getline` / 9-29
Printing output: `print` and `printf` / 9-30
`print` / 9-31
`printf` / 9-33
The `system` command / 9-34
Directing output to other programs / 9-34
Data structures / 9-35
Variables / 9-35
Initialization of variables / 9-37
Assignment operators / 9-37
Arrays / 9-38
Built-in variables and arrays / 9-40
Expressions / 9-41
Combining true-or-false expressions / 9-45
Implied concatenation operations / 9-45
Determination of data type / 9-46
Built-in string functions / 9-47
Built-in numeric functions / 9-49
Lexical conventions / 9-50
Numeric constants / 9-50
String constants / 9-51
Predefined variables, reserved keywords, and reserved function names / 9-51
Identifiers / 9-52
Record and field tokens / 9-52
Separators / 9-53
Record separators / 9-53
Field separator / 9-53
Multiline records / 9-54
Output record and field separators / 9-54
Separators and braces / 9-54
Primary expressions / 9-55
Numeric constants / 9-55
String constants / 9-56
Variables / 9-56
Functions / 9-57

Terms / 9-58
Binary terms / 9-58
Unary terms / 9-59
Incremented variables / 9-59
Terms with parentheses / 9-60
Expressions / 9-60
Concatenation of terms / 9-60
Assignment expressions / 9-61

Part 3 Math Tools

10 dc: A Desk Calculator / 10-1

Using dc / 10-2
Command syntax / 10-2
Operators / 10-3
Relational operators / 10-3
dc command set / 10-4
Input/output format and base / 10-4
Input conversion and base / 10-4
Output commands / 10-5
Scale / 10-5
Stack commands / 10-6
Subroutine definitions and calls / 10-6
Internal registers / 10-6
Pushdown registers and arrays / 10-7
Miscellaneous commands / 10-7
dc command quick reference / 10-8
Programming dc / 10-9

11 bc: A Basic Calculator / 11-1

Using bc / 11-3
bc command syntax / 11-3
Entering a program at the terminal / 11-4
Program files / 11-4
Exiting from bc / 11-4

Program syntax / 11-5
Comments / 11-5
Constants / 11-6
Keywords / 11-6
Identifiers / 11-6
Defining functions / 11-7
Function calls and function arguments / 11-7
The `return` statement / 11-8
Automatic variables / 11-8
Global variables / 11-9
Arrays or subscripted variables / 11-9
Statements / 11-10
Assignment statements / 11-12
Control statements / 11-13
Relational operators / 11-13
The `if` statement / 11-14
The `while` statement / 11-14
The `for` statement / 11-15
Expressions / 11-15
Input and output bases: `ibase` and `obase` / 11-17
`ibase` / 11-17
`obase` / 11-18
`scale` / 11-19

Part 4 Screen-Oriented Tools

12 `curses`: Terminal-Independent Screen I/O / 12-1

Overview of `curses` usage / 12-3
Output / 12-4
Input / 12-5
Highlighting / 12-8
Multiple windows / 12-10
Multiple terminals / 12-11
Low-level `terminfo` usage / 12-13
A larger example / 12-16
List of `curses` routines / 12-18
Structure / 12-18
Initialization / 12-19
Option setting / 12-20

Terminal mode setting / 12-23
Window manipulation / 12-24
Causing output to the terminal / 12-25
Writing on window structures / 12-27
Moving the cursor / 12-27
Writing one character / 12-27
Writing a string / 12-28
Clearing areas of the screen / 12-28
Inserting and deleting text / 12-29
Formatted output / 12-30
Miscellaneous / 12-30
Input from a window / 12-30
Input from the terminal / 12-31
Video attributes / 12-32
Bells and flashing lights / 12-33
Portability functions / 12-33
Delays / 12-34
Lower-level functions / 12-35
Cursor motion / 12-35
`terminfo` level / 12-35
Operation details / 12-39
Insert and delete line and character / 12-39
Additional terminals / 12-40
Multiple terminals / 12-40
Video attributes / 12-41
Special keys / 12-42
Scrolling region / 12-43
`mini-curses` / 12-43
TTY-mode functions / 12-45
Typeahead check / 12-45
`getstr` / 12-46
`longname` / 12-46
`nodelay` mode / 12-46
Portability / 12-46
Example program: `scatter` / 12-47
Example program: `show` / 12-49
Example program: `highlight` / 12-51
Example program: `window` / 12-53
Example program: `two` / 12-55
Example program: `termhl` / 12-59
Example program: `editor` / 12-62

13 Commando / 13-1

Introduction / 13-2
Macintosh dialog boxes / 13-3
Commando dialog boxes / 13-4
The Commando script language / 13-5
Dialog box layout / 13-5
Layout examples / 13-8
Single-row example / 13-8
Multiple-row example / 13-10
Column example / 13-12
Nested dialog box example / 13-14
Control examples / 13-16
Checkbox / 13-16
Radio buttons / 13-17
Text boxes / 13-19
Text / 13-21
Buttons / 13-21
Dependencies / 13-25
Boxes / 13-28
Leniencies / 13-28
Keywords / 13-28
Creating Commando dialogs / 13-30
Invoking Commando dialogs / 13-30
Writing Commando dialogs / 13-31
Testing Commando dialogs / 13-31
Compiling Commando dialogs / 13-32
Dialog design guidelines / 13-32
Dialog layout guidelines / 13-32
Dialog aesthetics / 13-33
Descriptive information / 13-34

Index / In-1

Figures, Tables, and Listings

Chapter 3 `yacc`: A Compiler-Writing System

Table 3-1 C language escapes recognized by `yacc` / 3-7
Table 3-2 Arithmetic operators / 3-39

Chapter 4 `m4`: A Macro Processor

Table 4-1 Arithmetic operators / 4-9

Chapter 5 `lex`: A Lexical Analyzer

Figure 5-1 Overview of `lex` / 5-3
Figure 5-2 `lex` with `yacc` / 5-5

Table 5-1 Regular expression operators / 5-31

Chapter 7 `make`: A File Production Tool

Listing 7-1 Sample listing of default rules file / 7-23
Listing 7-2 Replacing a default rule / 7-28

Table 7-1 Default suffix list / 7-21
Table 7-2 Macro names and default compilers / 7-27

Chapter 8 SCCS Reference

Listing 8-1 Sample interface program for group projects / 8-10

Figure 8-1 A linear progression of versions / 8-14
Figure 8-2 A branching SCCS tree / 8-15
Figure 8-3 A complicated branch structure / 8-15
Figure 8-4 Relationships among temporary files / 8-18
Figure 8-5 Removing a delta / 8-41

Table 8-1 SCCS ID keywords / 8-20
Table 8-2 Determination of a new SID / 8-35

Chapter 9 `awk` Programming Language

Table 9-1 Arithmetic operators / 9-42
Table 9-2 Assignment operators / 9-42
Table 9-3 Relational operators / 9-43
Table 9-4 Logical operators / 9-44
Table 9-5 Regular expression pattern-matching operators / 9-44
Table 9-6 Reserved strings / 9-52
Table 9-7 Values for sample numeric constants / 9-55
Table 9-8 Values for sample string constants / 9-56

Chapter 10 `dc`: A Desk Calculator

Table 10-1 `dc` operators / 10-3

Chapter 11 `bc`: A Basic Calculator

Table 11-1 Assignment statements / 11-12
Table 11-2 Relational operators / 11-14
Table 11-3 Operators and their precedence / 11-16

Chapter 12 `curses`: Terminal-Independent Screen I/O

Listing 12-1 Framework of a `curses` program / 12-3
Listing 12-2 Sending a message to several terminals / 12-13
Listing 12-3 `terminfo`-level framework / 12-14

Chapter 13 Commando

Figure 13-1 Schematic dialog box / 13-3
Figure 13-2 Commando dialog box for the UNIX command `lpr` / 13-4
Figure 13-3 Commando dialog box for the UNIX command `tar` / 13-5
Figure 13-4 Dialog box layout example / 13-6
Figure 13-5 Single-row dialog box / 13-8
Figure 13-6 Multiple-row dialog box / 13-10
Figure 13-7 Multiple-column dialog box / 13-12
Figure 13-8 Further dialog example / 13-14
Figure 13-9 Checkbox dialog example / 13-16
Figure 13-10 Radio button dialog example / 13-17
Figure 13-11 Text box dialog example / 13-19
Figure 13-12 Button example: Initial dialog box / 13-22
Figure 13-13 Button example: Save a File dialog box / 13-22
Figure 13-14 Button example: Save a File control was selected / 13-22
Figure 13-15 Button example: Redirection subdialog box / 13-23
Figure 13-16 Dependencies example: First control selected / 13-25
Figure 13-17 Dependencies example: Second control selected / 13-26

Listing 13-1 Dialog box layout example script / 13-7
Listing 13-2 Single-row dialog script / 13-8
Listing 13-3 Multiple-row dialog script / 13-11
Listing 13-4 Multiple-column dialog script / 13-12
Listing 13-5 Further dialog script / 13-15
Listing 13-6 Checkbox example script / 13-16
Listing 13-7 Radio button example script / 13-18
Listing 13-8 Text box example script / 13-20

Listing 13-9 Button example script / 13-24
Listing 13-10 Dependencies example script / 13-27

Table 13-1 File dialog keywords / 13-23
Table 13-2 Commando keyword reference / 13-29

About This Guide

This guide describes many A/UX tools to assist in program management and other tasks. This guide details program development tools to improve program structure, monitor program executions, and debug programs. Tools to assist in file management tasks, such as finding files, determining file characteristics, and maintaining groups of files, are also detailed in this guide. And finally, tools for processing and parsing text and code are described.

Who should use this guide

This guide is intended for programmers and developers. This guide does not serve as a tutorial to help you learn programming skills; rather, it serves as a reference to determine what tools are available in A/UX and how to use them effectively.

What you need to know

To get the most out of this guide, you need to have a good working knowledge of programming practices. This guide assumes that you are conversant with a programming language and with the general process of coding, compiling, testing, debugging, and so forth. A general knowledge of UNIX® is also assumed. You need to know the basic skills of using a Macintosh, such as double-clicking to open a file and dragging the mouse to choose a menu command.

What's covered in this guide

This guide describes the following topics:

- A/UX program development tools
- a compiler-writing system, `yacc`
- a macro processor, `m4`
- a lexical analyzer, `lex`
- file manipulation tools
- program maintenance tool, `make`
- version management tools for source code, SCCS
- a file-processing language, `awk`
- desk calculators, `dc` and `bc`
- terminal-independent input and output, `curses`
- screen-oriented input and output through Macintosh dialog boxes, Commando

If you need information about the tools directly involved in the compilation process, such as compilers (`cc` and `f77`), assemblers (`as`), link-editors (`ld`), and debuggers (`sdb` and `dbx`) see *A/UX Programming Languages and Tools*, Volume 1. Volume 1 also covers various libraries, the `lint` tool, `efl`, and the POSIX environment.

Where to go for more information

If you need information about the tools directly involved in the compilation process, see *A/UX Programming Languages and Tools*, Volume 1. If you need more information about the Macintosh interface, see *A/UX Toolbox: Macintosh ROM Interface*. If you would like information about porting applications to A/UX, see the *A/UX Porting Guide*.

How to use this guide

This guide serves as a reference to help you when programming and using these tools. As a reference book, it is not designed to be read from cover to cover. Each chapter is a discrete description of a particular tool or class of tools; therefore, you should skip directly to these compact references.

Conventions used in this guide

A/UX guides follow specific conventions. For example, words that require special emphasis appear in specific fonts or font styles. The following sections describe the conventions used in all A/UX guides.

Keys and key combinations

Certain keys on the keyboard have special names. These modifier and character keys, often used in combination with other keys, perform various functions. In this guide, the names of these keys are in Initial Capital letters followed by SMALL CAPITAL letters.

The key names are

CAPS LOCK	DOWN ARROW (↓)	OPTION	SPACE BAR
COMMAND (⌘)	ENTER	RETURN	TAB
CONTROL	ESCAPE	RIGHT ARROW (→)	UP ARROW (↑)
DELETE	LEFT ARROW (←)	SHIFT	

Sometimes you will see two or more names joined by hyphens. The hyphens indicate that you use two or more keys together to perform a specific function. For example,

Press COMMAND-K

means "Hold down the COMMAND key and press the K key."

Terminology

In A/UX guides, a certain term can represent a specific set of actions. For example, the word *enter* indicates that you type a series of characters on the command line and press the RETURN key. The instruction

Enter `ls`

means "Type `ls` and press the RETURN key."

Here is a list of common terms and the corresponding actions you take:

Term	*Action*
Click	Press and then immediately release the mouse button.
Drag	Position the mouse pointer, press and hold down the mouse button while moving the mouse, and then release the mouse button.
Choose	Activate a command in a menu. To choose a command from a pull-down menu, click once on the menu title and, while holding down the mouse button, drag down until the command is highlighted. Then release the mouse button.
Select	Highlight a selectable object by positioning the mouse pointer on the object and clicking.
Type	Type an entry *without* pressing the RETURN key.
Enter	Type the series of characters indicated and press the RETURN key.

The `Courier` font

Throughout A/UX guides, words that you see on the screen or that you must type exactly as shown are in the `Courier` font. For example, suppose you see this instruction:

Type `date` on the command line and press RETURN.

The word `date` is in the `Courier` font to indicate that you must type it. Suppose you then read this explanation:

Once you press RETURN, you'll see something like this:

```
Tues Oct 17 17:04:00 PDT 1989
```

In this case, `Courier` is used to represent exactly what appears on the screen.

All A/UX manual page names also are shown in the `Courier` font. For example, the entry `ls`(1) indicates that `ls` is the name of a manual page in an A/UX reference manual. See "Manual Page Reference Notation" below for more information on A/UX command reference manuals.

Font styles

Italics are used to indicate that a word or set of words is a placeholder for part of a command. For example,

`cat` *filename*

tells you that *filename* is a placeholder for the name of a file you wish to view. If you want to view the contents of a file named `Elvis`, type the word `Elvis` in place of *filename*. In other words, enter

```
cat Elvis
```

New terms appear in **boldface** where they are defined.

A/UX command syntax

A/UX commands follow a specific command syntax. A typical A/UX command gives the command name first, followed by options and arguments. For example, here is the syntax for the `wc` command:

`wc [-l] [-w] [`*directory*`]`...

In this example, `wc` is the command, `-l` and `-w` are options, *directory* is an argument, and the ellipses (...) indicate that more than one argument can be used. Note that each command element is separated by a space.

The following list gives more information about the elements of an A/UX command:

Element	*Description*
command	The command name.
option	A character or group of characters that modifies the command. Most options have the form -*option,* where *option* is a letter representing an option. Most commands have one or more options.
argument	A modification or specification of a command, usually a filename or symbols representing one or more filenames.

[ ] Brackets used to enclose an optional item—that is, an item that is not essential for execution of the command.

... Ellipses used to indicate that more than one argument can be entered.

For example, the `wc` command is used to count lines, words, and characters in a file. Thus, you can enter

```
wc -w Priscilla
```

In this command line, `-w` is the option that instructs the command to count all of the words in the file, and the argument `Priscilla` is the file to be searched.

Manual page reference notation

A/UX Command Reference, A/UX Programmer's Reference, A/UX System Administrator's Reference, X11 Command Reference for A/UX, and *X11 Programmer's Reference for A/UX* contain descriptions of commands, subroutines, and other related information. Such descriptions are known as *manual pages* (often shortened to *man pages*). Manual pages are organized within these references by section numbers. The standard A/UX cross-reference notation is

command (section)

where *command* is the name of the command, file, or other facility; *section* is the number of the section in which the item resides.

- Items followed by section numbers (1M) and (8) are described in *A/UX System Administrator's Reference.*
- Items followed by section numbers (1) and (6) are described in *A/UX Command Reference.*
- Items followed by section numbers (2), (3), (4), and (5) are described in *A/UX Programmer's Reference.*
- Items followed by section number (1X) are described in *X11 Command Reference for A/UX.*
- Items followed by section numbers (3X) and (3Xt) are described in *X11 Programmer's Reference for A/UX.*

For example,

```
cat(1)
```

refers to the command `cat`, which is described in Section 1 of *A/UX Command Reference.*

You can display manual pages on the screen by using the `man` command. For example, enter the command

```
man cat
```

to display the manual page for the `cat` command, including its description, syntax, options, and other pertinent information. To exit, press the SPACE BAR until you see a command prompt, or type `q` at any time to return immediately to your command prompt.

For more information

To find out where you need to go for more information about how to use A/UX, see *Road Map to A/UX.* This guide contains descriptions of each A/UX guide and ordering information for all the guides in the A/UX documentation suite.

Part 1 Program Development Tools

This section describes several tools you might find useful during program development and execution. The chapter "Programming Tools" describes utilities for

- structuring programs: `cb`
- observing program execution: `cflow` and `prof`
- processing: `cpp`
- finding functions in programs: `ctags`
- sharing strings in C programs: `xstr`
- debugging: `nm` and `od`

The following chapters in this section describe tools for

- a macro processor: `m4`
- a lexical analyzer: `lex`
- a compiler-writing system: `yacc`

1 Overview of Programming Tools

Program development tools / 1-2
File manipulation tools / 1-4
Math functions: `dc` and `bc` / 1-6
Screen-oriented tools: `curses` and Commando / 1-6

The A/UX environment provides many varied and useful tools to assist in program development and other related tasks. Tools are provided to find file characteristics, parse and process files, perform math, and control functions on the screen. This chapter provides a brief description of many of the tools and what primary tasks each one performs. The remaining chapters provide a more detailed discussion of these tools.

For information about tools directly related to the compilation process—the compilers, the assembler, the link editors, libraries, and debuggers—see *A/UX Programming Languages and Tools*, Volume 1.

Program development tools

In addition to the tools used for program compilation discussed in *A/UX Programming Languages and Tools*, Volume 1, A/UX offers several tools related to program development. These tools perform a variety of functions ranging from improving the format of your code to tracing your program execution and providing additional information for debugging. This section outlines these tools.

Program structure: `cb`

You can use the `cb` utility to improve the legibility and structure of C code. The `cb` utility reads C programs and writes them to the standard output with spacing and indentation that display the structure of the code.

Execution: `cflow` and `prof`

A/UX provides several tools for tracking the execution of a program. You can create a C flowgraph for a program using `cflow`. A C flowgraph shows how the program is put together, the flow control of the program, and how the subroutines are called. This flowgraph shows the order in which routines are called *graphically,* by level of indentation. The graph is built of external references, which include globals and function calls.

Another utility to show program execution is `prof`, which displays profile data on the running of a program to aid in optimization of the program. For each function, it gives the percentage of time spent executing it, the number of times it was called, and the time (in milliseconds) per call. You must compile your program with a special option to enable this capability.

Processing: `m4`, `lex`, and `yacc`

The A/UX environment includes several tools for processing text and code. This section provides a brief description of some of the more useful tools.

If you need a macro facility, you can use `m4` instead of `cpp`. `m4` is a general-purpose macro processor. The primary function of `m4` is to allow the replacement of certain text by other text. The `m4` utility reads every alphanumeric token (string of letters and digits) in the input and determines whether the token is the name of a macro. It then replaces the names of a macros by their defining text and pushes the resulting strings back into the input to be rescanned.

In addition to the straightforward replacement of one string of text by another, the `m4` macro processor also provides arguments to macros, arithmetic capabilities, file manipulation, conditional macro expansion, string and substring functions, and recursive definitions.

Another type of processor is `lex`. It is designed for lexical processing of character input streams. `lex` accepts high-level, problem-oriented specifications for character string matching. The `lex` utility can be useful when writing programs involving regular expressions as input and formatting input for parsing.

The `yacc` program is a parser generator used to impose structure on program input. After you create a specification of the input process, `yacc` generates a parser function, which calls the user-supplied low-level input routine (the lexical analyzer) to pick up the basic items, called "tokens," from the input stream. Tokens are organized according to the input structure rules, called "grammar rules." When one of these rules has been recognized, the user code (the "action") supplied for this rule is invoked. Actions have the ability to return values and make use of the values of other actions.

Debugging: `nm` and `od`

This book outlines a few tools useful in the debugging stage of program development. (The primary A/UX debuggers are detailed in *A/UX Programming Languages and Tools,* Volume 1.) The `nm` utility writes the symbol table for a Common Object File Format (COFF) file to standard output. `nm` lists each symbol and its value along with the location at which it is stored in memory.

The `od` command (octal dump) provides a means for examining binary files (usually unreadable on A/UX systems). If you need to know the function and procedure of some file available only in binary, you can use the `od` command with various options to discover what the file contains. The options correspond to available formats for interpreting bytes, characters, or words. If no options are specified, you can obtain a true octal dump, as words are interpreted in octal.

File manipulation tools

The A/UX tools detailed in this section help you perform file-related tasks such as finding a file size or location, determining the differences between two files, and obtaining the version of a program. Additionally, A/UX provides tools to control the file versions to ensure that they are the most recent and provides a way of updating and maintaining groups of files. The final tools in this section help you maintain current library archives and provide you with a file-processing language for parsing files.

File characteristics: `size`, `diff`, and `comm`

Often, you need to know characteristics of files. Some of the tools needed to obtain these attributes are briefly discussed here.

The `size` command produces size information for each section in the common object format files. The name of the section is shown followed by its size in bytes, physical address, and virtual address.

A/UX includes a number of programs that compare files to find differences, including `diff`, `bdiff`, `diff3`, `diffmk`, `diffdir`, `sdiff`, `cmp`, and `comm`. These programs all compare files or directories for differences.

The `find` command helps you locate files based on certain characteristics such as name, group, owner name, time of last modification or access, and so on. This powerful utility performs a recursive search for files of the given characteristics.

Maintenance: `make`, SCCS, and `ar`

The A/UX environment includes tools to update and maintain groups of files and to control the accessible versions of files to ensure that they are the most recent. Commands also exist to obtain the version number of programs you are running and to maintain up-to-date library archives.

The `make` program is a program-maintenance tool that keeps track of (and updates) groups of related files. All information about special libraries, special treatments, or options necessary for compiling multiple files is contained in a `make` description file. Using it ensures that your compilations reflect your latest changes.

The source code control system (SCCS) and revision control system (RCS) are version-management tools for source code or text files. In group projects, SCCS and RCS prevent multiple inconsistent versions of files from accumulating in several places. A single user can store multiple versions of a file without using a lot of disk space, easily reconstruct previous versions of a file, and keep track of versions with a simple, consistent numbering scheme.

The `version` command is useful for determining which version of a program you are running. `version` takes a list of files and reports the version number for each. The `version` command also reports the object file format of each file; that is, either `Coff object file format`, or `Old a.out object file format`.

You can use the archive command `ar` to combine several files into one archive. Archives consist of a collection of files, plus a table of contents. They are used mainly as libraries to be searched by the link editor `ld`. A library (or library archive) is an archive that contains object files (plus a table of contents). Putting together your own library allows you to use locally produced functions (instead of limiting you to the functions supplied in standard libraries). `ar` also provides the facility to append and delete archive files. Putting together your own library allows you to use locally produced functions (instead of limiting you to the functions supplied in standard libraries). With the `ar` command you can also move files around within the archive, as well as extract them, print them, and produce a table of contents.

A file-processing language: `awk`

The `awk` programming language is a file-processing language designed to make common information retrieval and manipulation tasks easy. The `awk` language can be used to generate reports, match patterns, validate data, or filter data for transmission.

Math functions: `dc` and `bc`

A/UX provides two specialized tools for handling arbitrary precision arithmetic, `dc` and `bc`. The `dc` program is an interactive desk calculator program. It has provisions for manipulating scaled fixed-point numbers and for input and output in bases other than decimal. `bc` is a specialized language and compiler for handling arbitrary precision arithmetic using the `dc` program.

Screen-oriented tools: `curses` and Commando

A/UX also provides the `curses` package to write screen-oriented programs. `curses` provides a terminal-independent method of screen-oriented input and output. It includes facilities for taking input from the terminal, sending output to a terminal, creating and manipulating windows on the screen, and performing screen updates in an optimal fashion. A program using the `curses` routines and functions generally needs to know nothing about the capabilities of any particular terminal; these characteristics are determined at execution time and guide the program in taking input and producing output. Thus, programs using this package can interact with a large variety of terminals and terminal types.

The Commando tool is useful for screen-oriented input and output on Macintosh computers. Commando lets you create CommandShell command lines by selecting controls within Macintosh dialog boxes. Controls direct the placement of options on the command line. When the user selects a particular control, Commando places a specific option on the command line. Once they are constructed, the command lines are either placed in a CommandShell window for execution or executed in a subshell.

2 Programming Tools

Improving C program structure: `cb` / 2-2
Generating a C flowgraph: `cflow` / 2-2
Displaying profile data: `prof` / 2-2
A C language preprocessor: `cpp` / 2-3
Finding a function definition quickly: `ctags` / 2-4
Sharing strings from C programs: `xstr` / 2-4
Printing the symbol table for a COFF file: `nm` / 2-5
Obtaining an octal dump of a file: `od` / 2-5

A/UX offers several tools related to program development. These tools perform a variety of functions, ranging from improving the format of your code to tracing your program execution and providing additional information for debugging. This chapter outlines many of these tools. The primary tools used for program compilation (the compilers, assembler, link editor, debuggers, and libraries) are discussed in *A/UX Programming Languages and Tools,* Volume 1.

Improving C program structure: `cb`

`cb` is used to improve the legibility and structure of C code. It reads C programs either from its arguments or from the standard input and writes them on the standard output with spacing and indentation that display the structure of the code. See `cb`(1) in *A/UX Command Reference* for more information.

Generating a C flowgraph: `cflow`

`cflow` generates a C flowgraph. A C flowgraph gives an idea of the following program features:

- how the program is put together
- the program flow of control
- how subroutines are called (that is, by which other routines and in which order)

This flowgraph shows the order in which routines are called *graphically*, by level of indentation. The graph is built of external references, which include globals and function calls. See `cflow`(1) in *A/UX Command Reference* for more information.

Displaying profile data: `prof`

`prof` displays profile data on the running of a program to aid in its optimization. For each function or global, it gives the percentage of time spent executing it, the number of times it was called, and the time (in milliseconds) per call. You must compile your program with a special option to enable profiling (see `cc`(1) in *A/UX Command Reference* for more details). See `prof`(1) in *A/UX Command Reference* for more information.

A C language preprocessor: `cpp`

You can use `cpp`, the C preprocessor, as a simple programming language that takes less time to compile than more complex languages. It strips comments, expands macros into their definitions, allows files to be read in (through `#include` statements), and provides a facility for conditional command execution. This means that you can intersperse text with comments. Comments are stripped; commands are executed.

Normally, `cpp` is invoked automatically as (the first) part of the `cc` command.

If you need a macro facility, you can use `m4` instead of `cpp`. `m4` is generally much more powerful than `cpp` as a macro processor. (For instance, `m4` allows recursive macro substitutions, while `cpp` does not.)

`cpp` is useful for

- stripping comments
- standardizing included definitions among many files for one project
- debugging (certain commands executed if in this mode, others if not)
- minimizing file space, combining many files into one

One of the most useful applications of `cpp` is as a debugging and program-control tool. Any statement included after an `#ifdef` *definition* is executed only if the *definition* was actually defined previously by means of a `#define` statement (or a `-D`*definition* in the command line). If not, and if there is an `#else` present, the statements between it and the `#endif` are executed. Otherwise, control is resumed at the level of the statement immediately following `#endif`. See `cpp`(1) in *A/UX Command Reference* for more information.

Finding a function definition quickly: `ctags`

Programs can rapidly accumulate a large number of functions, either in one source file or scattered across many files. `ctags` goes through the files given as its arguments and creates a new file, called `tags`. Each line in the file `tags` contains the following components:

- the name of one function
- where that function is located
- a scanning pattern that can be used to find the function

Unless `ctags` is used with either the `-a` (append) or the `-u` (update) option, a new `tags` file is created each time it is invoked.

Once the `tags` file is created, it can be accessed (thanks to the scanning pattern in the last field of each line) from `vi` (also from `ex`) by typing

`:ta` *function-name*

This causes the named function to appear on the editor's screen.

`ctags` can be used on Fortran and Pascal sources as well as C programs. See `ctags`(1) in *A/UX Command Reference.*

Sharing strings from C programs: `xstr`

The object of using `xstr` is to share one copy of a string among several files. If you need to modify the string throughout your program, you can modify it once instead of doing global searches through all your modules. If you have, in two different files,

```
char *ptr1 = "blah";
char *ptr2 = "blah";
```

`xstr` combines this into one string, in its `strings` file, and replaces occurrences of the string in the original files with a pointer to this string. This allows for shared constant strings among several files, or possibly among several users.

In practice, use of `xstr` can save memory space. After making the `xstr` array read only, you can arrange to have multiple users share these strings, thereby saving even more memory space. See `xstr`(1) in *A/UX Command Reference* for more information.

Printing the symbol table for a COFF file: `nm`

`nm` writes the symbol table for a COFF file to standard output. This is useful for debugging. `nm` lists each symbol and its value, along with the location at which it is stored in memory. See `nm`(1) in *A/UX Command Reference* for more information.

Obtaining an octal dump of a file: `od`

`od` provides a means for examining binary files (usually unreadable on A/UX systems). If you need to know the function and procedure of some file available only in binary, you can try the `od` command with various options to discover what the file contains. The options correspond to available formats for interpreting bytes, characters, or words. If no options are specified, a true octal dump is obtained, as words are interpreted in octal. See `od`(1) in *A/UX Command Reference* for more information.

You can also use the `strings` program to write the printable ASCII strings in a binary file onto standard output. This is useful for identifying unknown binary files. See `strings`(1) in *A/UX Command Reference* for more information.

3 `yacc`: A Compiler-Writing System

Usage / 3-3

Basic specifications / 3-6

Actions / 3-8

Lexical analysis / 3-12

Parser operation / 3-14

Ambiguity and conflicts / 3-19

Precedence / 3-23

Error handling / 3-27

The `yacc` environment / 3-29

Input style / 3-31

Left recursion / 3-31

Lexical considerations / 3-32

Reserved words / 3-34

Simulating error and accept in actions / 3-34

Accessing values in enclosing rules / 3-34

Arbitrary value types / 3-36

Example: A desk calculator / 3-38
Example: `yacc` input syntax / 3-43
Example: An advanced grammar / 3-46
Backward compatibility / 3-56

The `yacc` program is a general tool for imposing structure on the input to a computer program. `yacc` converts context-free grammar into a set of tables for a simple automaton that executes an `lr`(1) parsing algorithm. The grammar can be ambiguous; specified precedence rules are used to break ambiguities.

Usage

The first step in using `yacc` is to create a specification of the input process, which includes rules describing the input structure, code to be invoked when these rules are recognized, and a low-level routine to do the basic input. `yacc` then generates a function to control the input process. This function, called a "parser," calls the user-supplied low-level input routine (the lexical analyzer) to pick up the basic items (called "tokens") from the input stream.

Tokens are organized according to the input structure rules called "grammar rules." When one of these rules is recognized, the user code supplied for this rule (that is, an action) is invoked. Actions have the ability to return values and make use of the values of other actions.

`yacc` is written in a portable dialect of the C language, and the actions and output subroutine are written in the C language as well. Moreover, many of the syntactic conventions of `yacc` follow those of the C language.

The heart of the input specification is a collection of grammar rules. Each rule describes an allowable structure and gives it a name. For example, one grammar rule might be

```
date : month_name day ',' year;
```

where `date`, `month_name`, `day`, and `year` represent structures of interest in the input process; presumably, `month_name`, `day`, and `year` are defined elsewhere. The comma (`,`) is enclosed in single quotes. This implies that the comma is to appear literally in the input. The colon and semicolon serve merely as punctuation in the rule and have no significance in controlling the input. With proper definitions, the following input might be matched by the rule given above:

```
July 4, 1776
```

An important part of the input process is carried out by the lexical analyzer. This user routine reads the input stream, recognizes the lower-level structures, and communicates these tokens to the parser. For historical reasons, a structure recognized by the lexical analyzer is called a *terminal symbol*, while the structure recognized by the parser is called a *nonterminal symbol*. To avoid confusion, terminal symbols are usually referred to as *tokens*.

There is considerable leeway in deciding whether to recognize structures using the lexical analyzer or grammar rules. For example, the following rules might be used in the preceding example:

```
month_name : 'J'  'a'  'n' ;
month_name : 'F'  'e'  'b'  ;
...
month_name : 'D'  'e'  'c'  ;
```

The lexical analyzer needs to recognize only individual letters, and `month_name` is a nonterminal symbol. Such low-level rules tend to waste time and space and might complicate the specification beyond the ability of `yacc` to deal with it. Usually, the lexical analyzer recognizes the month names and returns an indication that a `month_name` is seen. In this case, `month_name` is a token. Literal characters (such as the comma above) must also be passed through the lexical analyzer and are also considered tokens.

Specification files are very flexible. If the rule

```
date : month '/'  day '/'  year;
```

were added to the above example, entering `7/4/1776` would be equivalent to `July 4, 1776` on input. In most cases, this new rule could be "slipped in" to a working system with minimal effort and little danger of disrupting existing input.

The input being read might not conform to the specifications. These input errors are detected as early as is theoretically possible with a left-to-right scan. Thus, not only is the chance of reading and computing with bad input data substantially reduced, but the bad data can usually be quickly found. Error handling, provided as part of the input specifications, permits the reentry of bad data or the continuation of the input process after skipping over the bad data.

In some cases, `yacc` fails to produce a parser when given a set of specifications. For example, the specifications might be self-contradictory, or they might require a more powerful recognition mechanism than that available to `yacc`. The former cases represent design errors; the latter cases can often be corrected by making the lexical analyzer more powerful or by rewriting some of the grammar rules.

While `yacc` cannot handle all possible specifications, its power compares favorably with similar systems. Moreover, the constructions that are difficult for `yacc` to handle are also frequently difficult for human beings to handle. Some users have reported that the discipline of formulating valid `yacc` specifications for their input revealed errors of conception or design early in the program development.

`yacc` has been used extensively in numerous practical applications on the A/UX system, including the syntax checker `lint`, the Portable C Compiler, and a system for typesetting mathematics.

The remainder of this chapter describes

- basic process of preparing a `yacc` specification
- parser operation
- handling ambiguities
- handling operator precedence in arithmetic expressions
- error detection and recovery
- the operating environment and special features of the parsers `yacc` produces
- suggestions to improve the style and efficiency of the specifications
- advanced topics

In addition, there are four sections that illustrate the earlier material:

- "A Desk Calculator" contains a brief example of using `yacc` to design a simple program.
- "`yacc` Input Syntax" contains a summary of the `yacc` input syntax.
- "An Advanced Grammar" contains an example using some of the more advanced features of `yacc`.
- "Backward Compatibility" contains a description of the mechanisms and syntax that, though no longer actively supported, are provided for historical continuity with older versions of `yacc`.

Basic specifications

Names refer to either tokens or nonterminal symbols. `yacc` requires token names to be declared as such. In addition, it is often desirable to include the lexical analyzer as part of the specification file. It might be useful to include other programs as well.

Every specification file consists of three sections:

- declarations
- grammar rules
- programs

These sections are separated by double percent symbols (`%%`). The percent symbol is generally used in `yacc` specifications as an escape character.

The following is a syntactic description of a `yacc` specification file:

declarations

`%%`

rules

`%%`

programs

The *declarations* section might be empty and, if the *programs* section is omitted, the second `%%` mark might also be excluded. The smallest legal `yacc` specification is therefore

`%%`

rules

Blanks, tabs, and newlines are ignored, but they cannot appear in names or multicharacter reserved symbols. Comments can appear wherever a name is legal. They are enclosed in `/*` and `*/`, as in the C language.

The *rules* section is made up of one or more grammar rules. A grammar rule has the following form:

a : *body*;

In this example, *a* represents a nonterminal name, and *body* represents a sequence of zero or more names and literals. The colon and the semicolon are `yacc` punctuation.

Names can be of arbitrary length and can be made up of letters, dots, underscores, and noninitial digits. Uppercase and lowercase letters are distinct. The names used in the body of a grammar rule can represent tokens or nonterminal symbols.

A literal consists of a character enclosed in single quotes (`' '`).

As in the C language, the backslash (`\`) is an escape character within literals, and all the C language escapes are recognized. Table 3-1 lists the escapes recognized by `yacc`.

Table 3-1 C language escapes recognized by `yacc`

Escape	Meaning
`\n`	Newline
`\r`	Return
`\'`	Single quote (`'`)
`\\`	Backslash (`\`)
`\t`	Tab
`\b`	Backspace
`\f`	Form feed
`\`*xxx*	*xxx* in octal

For a number of technical reasons, the null character (`\0` or 0) should never be used in grammar rules.

If there are several grammar rules with the same left side, the vertical bar (`|`) can be used to avoid rewriting the left side. The semicolon at the end of a rule can be dropped before a vertical bar. Thus, the grammar rules

```
A : B C D;
A : E F;
A : G;
```

can be given to `yacc` using the vertical bar:

```
A : B C D
  | E F
  | G;
```

It is not necessary that all grammar rules with the same left side appear together in the grammar rules section, although it makes the input much easier to read and change.

If a nonterminal symbol matches the empty string, this can be indicated by

```
empty : ;
```

Names representing tokens must be declared in the declarations section. For example,

`%token` *name1 name2*

Every name not defined in the declarations section is assumed to represent a nonterminal symbol. Nonterminal symbols must appear on the left side of at least one rule.

The parser is designed to recognize the nonterminal start symbol. Thus, this symbol represents the largest, most general structure described by the grammar rules. By default, the start symbol is taken to be the left side of the first grammar rule in the *rules* section.

It is possible and desirable to declare the start symbol explicitly in the *declarations* section using the `%start` keyword. For example,

`%start` *symbol*

The end of the input to the parser is signaled by a special token, called the *end-marker*. If the tokens up to but not including the end-marker form a structure that matches the start symbol, the parser function returns to its caller after the end-marker is seen and accepts the input. If the end-marker is seen in any other context, it is an error.

It is the job of the user-supplied lexical analyzer to return the end-marker when appropriate. Usually the end-marker represents some reasonably obvious I/O status, such as end-of-file or end-of-record.

Actions

With each grammar rule, the user can associate actions to be performed each time the rule is recognized in the input process. These actions can return values and can obtain the values returned by previous actions. Moreover, the lexical analyzer can return values for tokens, if desired.

An action is an arbitrary C language statement and as such can do input and output, call subprograms, and alter external vectors and variables. An action is specified by one or more statements enclosed in braces (`{` and `}`).

For example,

```
A : '(' B ')'
{
 hello( 1,  "abc" );
}
```

and the following is an example of grammar rules with actions:

```
XXX : YYY ZZZ
    {
        printf("a message\n");
        flag = 25;
    }
```

To facilitate easy communication between the actions and the parser, the action statements are altered slightly. The dollar sign symbol (`$`) is used as a signal to `yacc` in this context. To return a value, the action normally sets the pseudovariable `$$` to some value. The following action does nothing except return the value of one:

```
{ $$ = 1; }
```

To obtain the values returned by previous actions and the lexical analyzer, the action can use the pseudovariables `$1`, `$2`, and so on, which refer to the values returned by the components of the right side of a rule, reading from left to right. For example, if the rule is

```
A : B C D;
```

then `$2` has the value returned by `C`, and `$3` the value returned by `D`.

With the following rule, the value returned is usually the value of the *expr* in parentheses:

```
expr : '(' expr ')'
    {
        $$ = $2 ;
    }
```

By default, the value of a rule is the value of the first element in it (`$1`).

Grammar rules of the following form frequently need not have an explicit action:

```
A : B;
```

In the preceding examples, all the actions came at the end of rules. Sometimes, though, it is desirable to obtain control before a rule is fully parsed. The `yacc` program permits an action to be written in the middle of a rule as well as at the end.

This kind of rule is assumed to return a value accessible through the usual `$` mechanism by the actions to the right of it. In turn, it can access the values returned by the symbols to the left of the action. For example, in the following rule `x` is set to 1 (the value returned by the action to its left) and `y` is set to the value returned by `C`:

```
A : B
    {
        $$ = 1;
    }
        C
    {
        x = $2;
        y = $3;
    }
;
```

This is because every component of the right side of the rule, including an action, is associated with a positional pseudovariable, so the `$1` refers to `B`, `$2` to the value returned by the action associated with `B`, `$3` to `C`, and so on.

Actions that do not terminate a rule are actually handled by `yacc` by manufacturing a new nonterminal symbol name and a new rule matching this name to the empty string. The interior action is the action triggered by recognizing this added rule.

`yacc` actually treats the preceding example as if it were written like the following example (`$ACT` is an empty action):

```
$ACT   :     /* empty */
    {
        $$ = 1;
    }
A   ;
      :  B $ACT C
```

```
{
    x = $2;
    y = $3;
}
;
```

In many applications, output is not produced directly by the actions. A data structure, such as a parse tree, is constructed in memory and transformations are applied to it before output is generated. Parse trees are particularly easy to construct, given routines to build and maintain the tree structure desired.

In the following example, the C function `node` creates a node with label *l* and descendants *n1* and *n2* and returns the index of the newly created node:

`node(`*l*`,` *n1*`,` *n2*`)`

Then a parse tree is built by supplying the actions following in the `yacc` specification file as follows:

```
expr : expr '+'  expr
    {
        $$ = node('+', $1, $3 );
    }
```

The user can define other variables to be used by the actions.

Declarations and definitions can appear in the declarations section enclosed in the marks `%{` and `%}`. These declarations and definitions have global scope, so they are known to the action statements and the lexical analyzer. For example,

```
%{ int variable = 0; %}
```

could be placed in the declarations section, making `variable` accessible to all of the actions.

The `yacc` parser uses only names beginning with `yy`. The user should avoid such names. In these examples, all the values are integers. A discussion of values of other types is found in the section "Arbitrary Value Types."

Lexical analysis

The user must supply a lexical analyzer to read the input stream and communicate tokens (with values, if desired) to the parser. The lexical analyzer is an integer-valued function called `yylex`. The function returns an integer, the *token number*, representing the kind of token read. If there is a value associated with that token, it should be assigned to the external variable `yylval`.

The parser and the lexical analyzer must agree on these token numbers for communication between them to take place. The numbers can be chosen by `yacc` or by the user. In either case, the `#define` mechanism of the C language is used to allow the lexical analyzer to return these numbers symbolically. For example, suppose that the token name `DIGIT` is defined in the declarations section of the `yacc` specification file. The relevant portion of the lexical analyzer might look like the following example:

```
yylex()
{
    extern int yylval;
    int c;
    ...
    c = getchar();
    ...
    switch( c )
    {
    ...
    case '0' :
    case '1' :
    ...
    case '9' :
        yylval = c - '0' ;
        return( DIGIT );
        ...
    }
    ...
}
```

The intent is to return a token number of `DIGIT` and a value equal to the numeric value of the digit. Provided that the lexical analyzer code is placed in the programs section of the specification file, the identifier `DIGIT` is defined as the token number associated with the token `DIGIT`.

This mechanism leads to clear, easily modified lexical analyzers. The only pitfall to avoid is using any token names in the grammar that are reserved or significant in the C language or the parser. For example, the use of token names `if` or `while` almost certainly causes severe difficulties when the lexical analyzer is compiled.

The token name `error` is reserved for error handling and should not be used naively.

As mentioned earlier, the token numbers can be chosen by `yacc` or by the user. In the default situation, the numbers are chosen by `yacc`. The default token number for a literal character is the numeric value of the character in the local character set. Other names are assigned token numbers starting at 257.

To assign a token number to a token (including literals), the first appearance of the token name or literal in the *declarations* section can be immediately followed by a non-negative integer. This integer is taken to be the token number of the name or literal. Names and literals not defined by this mechanism retain their default definitions. It is important that all token numbers be distinct.

For historical reasons, the end-marker must have token number 0 or be negative. This token number cannot be redefined by the user. Thus, all lexical analyzers should be prepared to return 0 or a negative number as a token upon reaching the end of their input.

The `lex` program is a very useful tool for constructing lexical analyzers. These lexical analyzers are designed to work in close harmony with `yacc` parsers. The specifications for these lexical analyzers use regular expressions instead of grammar rules.

`lex` can easily be used to produce quite complicated lexical analyzers, but there remain some languages (such as Fortran) that do not fit any theoretical framework and whose lexical analyzers must be crafted by hand. See Chapter 5 in this manual, "`lex`: A Lexical Analyzer," for more information on `lex`.

Parser operation

The `yacc` program turns the specification file into a C language program, which parses the input according to the specification given. The algorithm used to go from the specification to the parser is complex and is not discussed here. The parser itself, however, is relatively simple, and understanding how it works makes treatment of error recovery and ambiguities much more comprehensible.

The parser produced by `yacc` consists of a finite-state machine with a stack. The parser also is capable of reading and remembering the next input token (called the "look-ahead token"). The current state is always the one on the top of the stack. The states of the finite-state machine are given small integer labels.

Initially, the machine is in state 0 (the stack contains only state 0) and no look-ahead token has been read. The machine has only four actions available:

`shift`	Push current state onto stack; go into specified new state.
`reduce`	Pop some number of states from stack; push new state; execute user code.
`accept`	End of input has been (successfully) reached.
`error`	An unparsable situation has been detected.

A step of the parser is done as follows:

1. Based on its current state, the parser decides whether it needs a look-ahead token to choose the action to be taken. If it needs one and does not have one, it calls `yylex` to obtain the next token.
2. Using the current state and the look-ahead token if needed, the parser decides on its next action and carries it out. This can cause states to be pushed onto the stack or popped off the stack and the look-ahead token to be processed or left alone.

The `shift` action is the most common action the parser takes. Whenever a `shift` action is taken, there is always a look-ahead token. In the following example, in state 56, if the look-ahead token is `IF`, the current state (56) is pushed down on the stack, and state 34 becomes the current state (on the top of the stack):

```
IF shift 34
```

The look-ahead token is cleared.

The `reduce` action keeps the stack from growing without bounds. `reduce` actions are appropriate when the parser has seen the right side of a grammar rule and is prepared to announce that it has seen an instance of the rule replacing the right side by the left side.

It might be necessary to consult the look-ahead token to decide whether to reduce (usually it is not necessary). In fact, the default action (represented by a dot) is often a `reduce` action.

`reduce` actions are associated with individual grammar rules. Grammar rules are also given small integer numbers, and this leads to some confusion. For example, in the following display, the action refers to grammar rule 18:

```
. reduce 18
```

While in this example, the action refers to state 34:

```
IF shift 34
```

Suppose the following rule is being reduced:

```
A  : x y z ;
```

The `reduce` action depends on the left symbol (`A` in this case) and the number of symbols on the right side (three in this case). To reduce, first pop off the top three states from the stack. (In general, the number of states popped equals the number of symbols on the right side of the rule.) In effect, these states were the ones put on the stack while recognizing `x`, `y`, and `z`, and no longer serve any useful purpose.

After popping these states, a state is uncovered that was the state the parser was in before beginning to process the rule. Using this uncovered state and the symbol on the left side of the rule, perform what is, in effect, a shift of `A`. A new state is obtained and pushed onto the stack, and parsing continues.

There are significant differences between the processing of the left symbol and an ordinary shift of a token, however, so this action is called a `goto` action. In particular, the look-ahead token is cleared by a shift but is not affected by a `goto`. In any case, the uncovered state contains an entry such as the following one, which causes state 20 to be pushed onto the stack and become the current state:

```
A goto 20
```

In effect, the `reduce` action "turns back the clock" in the parse, popping the states off the stack to go back to the state where the right side of the rule was first seen. The parser then behaves as if it had seen the left side at that time. If the right side of the rule is empty, no states are popped off the stacks. The uncovered state is, in fact, the current state.

The `reduce` action also is important in the treatment of user-supplied actions and values. When a rule is reduced, the code supplied with the rule is executed before the stack is adjusted. In addition to the stack holding the states, another stack running in parallel with it holds the values returned from the lexical analyzer and the actions.

When a shift takes place, the external variable `yylval` is copied onto the value stack. After the return from the user code, the reduction is carried out. When the `goto` action is done, the external variable `yyval` is copied onto the value stack. The pseudovariables `$1`, `$2`, and so on refer to the value stack. The other two parser actions are conceptually much simpler. The `accept` action indicates that the entire input has been seen and that it matches the specification. This action appears only when the look-ahead token is the end-marker and indicates that the parser successfully did its job.

The `error` action, on the other hand, represents a place where the parser can no longer continue parsing according to the specification. The input tokens it has seen (together with the look-ahead token) cannot be followed by anything that would result in a legal input. The parser reports an error and attempts to recover the situation and resume parsing. The error recovery (as opposed to the detection of error) is discussed later.

Consider the following example as a `yacc` specification:

```
%token DING DONG DELL
%%
rhyme   :   sound place
        ;
sound   :   DING DONG
        ;
place   :   DELL
        ;
```

When `yacc` is invoked with the `-v` option, a file called `y.output` is produced with a human-readable description of the parser.

The following example is the `y.output` file corresponding to the above grammar (with some statistics stripped off the end), where the actions for each state are specified and there is a description of the parsing rules being processed in each state.

```
state 0
      $accept : _rhyme $end

      DING  shift 3
      .  error

      rhyme  goto 1
      sound  goto 2

state 1
      $accept :  rhyme_$end

      $end  accept
      . error

state 2
      rhyme :  sound_place

      DELL  shift 5
      . error

      place  goto 4

state 3
      sound :  DING_DONG

      DONG  shift 6
      . error

state 4
      rhyme :  sound place_     (1)

      . reduce  1

state 5
      place :  DELL_            (3)

      . reduce 3

state 6
      sound :  DING DONG_       (2)

      . reduce 2
```

The underscore character _ is used to indicate what was seen and what is yet to come in each rule.

The following input can be used to track the operations of the parser:

```
DING DONG DELL
```

Initially, the current state is state 0. The parser needs to refer to the input to decide between the actions available in state 0, so the first token (`DING`) is read and becomes the look-ahead token.

The action in state 0 on `DING` is `shift 3`. State 3 is pushed onto the stack, and the look-ahead token is cleared. State 3 becomes the current state. The next token (`DONG`) is read and becomes the look-ahead token. The action in state 3 on the token `DONG` is `shift 6`. State 6 is pushed onto the stack, and the look-ahead is cleared.

The stack now contains 0, 3, and 6. In state 6, without even consulting the look-ahead, the parser reduces by the following, which is rule 2:

```
sound :  DING DONG
```

Two states, 6 and 3, are popped off the stack, uncovering state 0. Consulting the description of state 0 (looking for a `goto` on `sound`), the following is obtained:

```
sound goto 2
```

State 2 is pushed onto the stack and becomes the current state. In state 2, the next token (`DELL`) must be read. The action is `shift 5`, so state 5 is pushed onto the stack, which now has 0, 2, and 5 on it, and the look-ahead token is cleared.

In state 5, the only action is to reduce by rule 3. This has one symbol on the right side, so one state (5) is popped off and state 2 is uncovered. The `goto` in state 2 on `place` (the left side of rule 3) is state 4. Now, the stack contains 0, 2, and 4.

In state 4, the only action is to reduce by rule 1. There are two symbols on the right, so the top two states are popped off, uncovering state 0 again. In state 0, there is a `goto` on `rhyme` causing the parser to enter state 1. In state 1, the input is read and the end-marker is obtained indicated by `$end` in the `y.output` file. The action in state 1 (when the end-marker is seen) successfully ends the parse.

The reader is urged to consider how the parser works when confronted with such incorrect strings as `DING DONG DONG`, `DING DONG`, `DING DONG DELL DELL`, and so on. A few minutes spent studying this and other simple examples can be repaid when problems arise in more complicated contexts.

Ambiguity and conflicts

A set of grammar rules is ambiguous if there is some input string that can be structured in two or more different ways. For example, the following grammar rule is a natural way of expressing the fact that one way of forming an arithmetic expression is to put two other expressions together with a minus sign between them:

expr : *expr* '-' *expr*

Unfortunately, this grammar rule does not completely specify the way that all complex inputs should be structured. For example, if the input is

expr - *expr* - *expr*

the rule allows this input to be structured as either

(*expr* - *expr*) - *expr*

or

expr - (*expr* - *expr*)

(The first is called *left association*, the second *right association*.) The `yacc` program detects such ambiguities when it is attempting to build the parser.

Consider the problem that confronts the parser when provided with the following input:

expr - *expr* - *expr*

When the parser has read the second *expr*, the input seen matches the right side of the previous grammar rule:

expr - *expr*

The parser can reduce the input by applying this rule. After applying the rule, the input is reduced to *expr* (the left side of the rule). The parser then reads the final part of the input (displayed in the following example) and again reduces:

- *expr*

The effect of this is to take the left associative interpretation. Alternatively, if the parser sees the following input:

expr - *expr*

it can defer the immediate application of the rule and continue reading the input until it sees the following input,

expr - *expr* - *expr*

It can then apply the rule to the right-most three symbols, reducing them to *expr*, which results in the following input being left:

expr - *expr*

Now the rule can be reduced once more. The effect is to take the right associative interpretation. The parser can do one of two legal things, a shift or a reduction. It has no way of deciding between them. This is called a *shift/reduce conflict.*

It might also happen that the parser has a choice of two legal reductions. This is called a *reduce/reduce conflict.* (Note that there are never any shift/shift conflicts.) When there are shift/reduce or reduce/reduce conflicts, `yacc` still produces a parser. It does this by selecting one of the valid steps wherever it has a choice.

A rule describing the choice to make in a given situation is called a *disambiguating rule.* The `yacc` program invokes two disambiguating rules by default:

- In a shift/reduce conflict, the default is to do the shift.
- In a reduce/reduce conflict, the default is to reduce by the earlier grammar rule (in the input sequence).

The first rule implies that reductions are deferred in favor of shifts when there is a choice. The second rule gives the user rather crude control over the behavior of the parser in this situation, but reduce/reduce conflicts should be avoided when possible.

Conflicts can arise because of mistakes in input or logic or because the grammar rules (while consistent) require a more complex parser than `yacc` can construct. The use of actions within rules can also cause conflicts if the action must be done before the parser can be sure which rule is being recognized. In these cases, the application of disambiguating rules is inappropriate and leads to an incorrect parser. For this reason, `yacc` always reports the number of shift/reduce and reduce/reduce conflicts resolved by rule 1 and rule 2.

In general, whenever it is possible to apply disambiguating rules to produce a correct parser, it is also possible to rewrite the grammar rules so that the same inputs are read but there are no conflicts. For this reason, most previous parser generators have considered conflicts to be fatal errors. Experience has suggested that this rewriting is somewhat unnatural and produces slower parsers. Thus, `yacc` produces parsers even in the presence of conflicts.

As an example of the power of disambiguating rules, consider

```
stat : IF '(' cond ')'  stat
     | IF '(' cond ')'  stat ELSE stat
     ;
```

which is a fragment from a programming language involving an if-then-else statement.

In these rules, `IF` and `ELSE` are tokens, `cond` is a nonterminal symbol describing conditional (logical) expressions, and `stat` is a nonterminal symbol describing statements. The first rule is called the *simple-if* rule and the second the *if-else* rule. These two rules form an ambiguous construction because input of the following form can be structured according to these rules in two ways:

```
IF ( C1 ) IF ( C2 ) S1 ELSE S2
```

The input can be structured as in the following example or as in the subsequent example, which is the one given in most programming languages having this construct:

```
IF ( C1 )
{
   IF ( C2 )
            S1
}
ELSE
     S2
or:
IF  ( C1 )
{
      IF  ( C2 )
            S1
     ELSE
            S2
}
```

Each `ELSE` is associated with the preceding `IF` that is without an `ELSE`.

In the following example, consider the situation where the parser has seen the `IF-ELSE` construct and is looking at the `ELSE`.

```
IF ( C1 ) IF ( C2 ) S1
```

It can immediately reduce by the simple-if rule to get

```
IF ( C1 ) stat
```

and then read the remaining input

```
ELSE S2
```

and reduce by the if-else rule. This leads to the first of the above groupings of the input.

On the other hand, the `ELSE` can be shifted, `S2` read, and then the right portion reduced by the if-else rule to get the following line, which can be reduced by the simple-if rule:

```
IF ( C1 ) stat
```

This leads to the second of the above groupings of the input, which is usually desired. Once again, the parser can do two valid things—there is a shift/reduce conflict. The application of disambiguating rule 1 tells the parser to shift in this case, which leads to the desired grouping. This shift/reduce conflict arises only when there is a particular current input symbol, `ELSE`, and particular inputs, such as have already been seen:

```
IF ( C1 ) IF ( C2 ) S1
```

In general, there might be many conflicts, and each one is associated with an input symbol and a set of previously read inputs. The previously read inputs are characterized by the "state" of the parser. The conflict messages of `yacc` are best understood by examining the verbose (`-v`) option output file. For example, the output corresponding to the above conflict state might be

```
23: shift/reduce conflict (shift 45, reduce 18) on ELSE
state 23
        stat   :  IF  (  cond  )  stat_          (18)
        stat   :  IF  (  cond  )  stat_ELSE  stat
        ELSE      shift 45
        .         reduce 18
```

where the first line describes the conflict, giving the "state" and the input symbol.

The ordinary state description gives the grammar rules active in the state and the parser actions.

Recall that the underline marks the portion of the grammar rules that has been seen. Thus, in the example, in state 23 the parser has seen input corresponding to `IF ( cond ) stat`, and the two grammar rules shown are active at this time.

The parser can do two things:

- If the input symbol is `ELSE`, it is possible to shift into state 45. State 45 has, as part of its description, the following line:

  ```
  stat  :  IF ( cond ) stat ELSE_stat
  ```

  because the `ELSE` will have been shifted in this state. In state 23, the alternative action (describing a dot (`.`)) is to be done if the input symbol is not mentioned explicitly in the actions.
- If the input symbol is not `ELSE`, the parser reduces to

  ```
  stat  : IF '(' cond ')'  stat
  ```

  by grammar rule 18.

Once again, notice that the numbers following `shift` commands refer to other states, while the numbers following `reduce` commands refer to grammar rule numbers.

In the `y.output` file, the rule numbers are printed after those rules that can be reduced. In most states, only one `reduce` action is possible, and it is the default command.

The user who encounters unexpected shift/reduce conflicts probably wants to look at the verbose output to decide whether the default actions are appropriate.

Precedence

There is one common situation where the rules given above for resolving conflicts are not sufficient. This is in the parsing of arithmetic expressions. Most of the commonly used constructions for arithmetic expressions can be naturally described by the notion of precedence levels for operators, together with information about left or right associativity.

It turns out that ambiguous grammars with appropriate disambiguating rules can be used to create parsers that are faster and easier to write than parsers constructed from unambiguous grammars. The basic notion is to write grammar rules of the following two forms for all binary and unary operators desired:

expr : *expr* `OP` *expr*

and

expr : `UNARY` *expr*

This creates a very ambiguous grammar with many parsing conflicts. As disambiguating rules, the user specifies the precedence or binding strength of all the operators and the associativity of the binary operators. This information is sufficient to allow `yacc` to resolve the parsing conflicts in accordance with these rules and construct a parser that realizes the desired precedences and associativities.

The precedences and associativities are attached to tokens in the *declarations* section. This is done by a series of lines beginning with one of the following `yacc` keywords: `%left`, `%right`, or `%nonassoc`, followed by a list of tokens. All of the tokens on the same line are assumed to have the same precedence level and associativity; the lines are listed in order of increasing precedence or binding strength. For example,

```
%left  '+'  '-'
%left  '*'  '/'
```

describes the precedence and associativity of the four arithmetic operators. Plus and minus are left associative and have lower precedence than star and slash, which are also left associative.

The keyword `%right` is used to describe right associative operators, and the keyword `%nonassoc` is used to describe operators like the operator `.LT` in Fortran that cannot associate with themselves. For example, the following line is illegal in Fortran and such an operator would be described with the keyword `%nonassoc` in `yacc`:

```
A  .LT.  B  .LT.  C
```

As an example of the behavior of these declarations, the following description might be used to structure the subsequent input:

```
%right  '='
%left  '+'  '-'
%left  '*'  '/'

%%
```

```
expr    :   expr    '='    expr
        |   expr    '+'    expr
        |   expr    '-'    expr
        |   expr    '*'    expr
        |   expr    '/'    expr
        |   NAME
        ;
```

The following line is the input to be structured by the above description to perform the correct precedence of operators:

```
a = b = c * d - e - f * g
```

The result of the structuring is as follows:

```
a = ( b = ( ((c*d)-e) - (f*g) ) )
```

When this mechanism is used, unary operators must, in general, be given a precedence. Sometimes a unary operator and a binary operator have the same symbolic representation but different precedences. An example is unary and binary minus (-). Unary minus can be given the same strength as multiplication, or even higher, while binary minus has a lower strength than multiplication.

The keyword `%prec` changes the precedence level associated with a particular grammar rule. `%prec` appears immediately after the body of the grammar rule, before the action or closing semicolon, and is followed by a token name or literal. The keyword causes the precedence of the grammar rule to become that of the following token name or literal. For example, the following rules might be used to give unary minus the same precedence as multiplication:

```
%left  '+'  '-'
%left  '*'  '/'

%%

expr    :   expr   '+'   expr
        |   expr   '-'   expr
        |   expr   '*'   expr
        |   expr   '/'   expr
        |   '-'    expr        %prec   '*'
        |   NAME
        ;
```

A token declared by `%left`, `%right`, and `%nonassoc` need not be, but can be declared by `%token` as well.

The precedences and associativities are used by `yacc` to resolve parsing conflicts. They give rise to disambiguating rules. Formally, the rules work as follows:

- The precedences and associativities are recorded for those tokens and literals that have them.
- A precedence and associativity is associated with each grammar rule. It is the precedence and associativity of the last token or literal in the body of the rule. If the `%prec` construction is used, it overrides this default. Some grammar rules can have no precedence and associativity associated with them.
- When there is a reduce/reduce conflict or there is a shift/reduce conflict and either the input symbol or the grammar rule has no precedence and associativity, then the two disambiguating rules given at the beginning of the section are used, and the conflicts are reported.
- If there is a shift/reduce conflict and both the grammar rule and the input character have precedence and associativity associated with them, then the conflict is resolved in favor of the action (shift or reduce) associated with the higher precedence. If the precedences are the same, then the associativity is used; left associative implies reduce, right associative implies shift, and nonassociating implies error.

Conflicts resolved by precedence are not counted in the number of shift/reduce and reduce/reduce conflicts reported by `yacc`. This means that mistakes in the specification of precedences can disguise errors in the input grammar. It is a good idea to be sparing with precedences and use them in an essentially "cookbook" fashion until some experience has been gained. The `y.output` file is very useful in deciding whether the parser is actually doing what was intended.

Error handling

Error handling is an extremely difficult area, and many of the problems are semantic ones. When an error is found, for example, it might be necessary to reclaim parse tree storage, delete or alter symbol table entries, and typically, set switches to avoid generating any further output. It is seldom acceptable to stop all processing when an error is found. It is more useful to continue scanning the input to find further syntax errors. This leads to the problem of getting the parser "restarted" after an error.

A general class of algorithms to do this involves discarding a number of tokens from the input string and attempting to adjust the parser so that input can continue. To allow the user some control over this process, `yacc` provides a simple but reasonably general feature. The token name `error` is reserved for error handling. This name can be used in grammar rules. In effect, it suggests places where errors are expected and recovery might take place.

The parser pops its stack until it enters a state where the token `error` is legal. It then behaves as if the token `error` were the current look-ahead token and performs the action encountered. The look-ahead token is then reset to the token that caused the error. If no special error rules are specified, the processing halts when an error is detected.

To prevent a cascade of error messages, the parser, after detecting an error, remains in error state until three tokens are successfully read and shifted. If an error is detected when the parser is already in error state, no message is given and the input token is quietly deleted.

As an example, a rule of the following form means that on a syntax error the parser attempts to skip over the statement in which the error is seen:

```
stat    :    error
```

More precisely, the parser scans ahead, looking for three tokens that might legally follow a statement, and starts processing at the first of these. If the beginnings of statements are not sufficiently distinctive, it might make a false start in the middle of a statement and end up reporting a second error where there is, in fact, no error.

Actions can be used with these special error rules. These actions might attempt to reinitialize tables, reclaim symbol table space, and so on. Error rules such as the ones mentioned are very general but difficult to control. Rules such as the following ones are somewhat easier. Here, when there is an error, the parser attempts to skip over the statement but does so by skipping to the next semicolon:

```
stat   :   error   ';'
```

All tokens after the error and before the next semicolon cannot be shifted and are discarded. When the semicolon is seen, this rule is reduced and any "cleanup" action associated with it performed.

Another form of error rule arises in interactive applications where it might be desirable to permit a line to be reentered after an error. The following example is one way to do this:

```
input   :  error    '\n'
            {
               printf("Reenter last line: ");
            }
               input
        {
            $$ = $4;
        }
        ;
```

There is one potential difficulty with this approach. The parser must correctly process three input tokens before it admits that it correctly resynchronized after the error. If the reentered line contains an error in the first two tokens, the parser deletes the offending tokens and gives no message. This is clearly unacceptable. For this reason, there is a mechanism that can force the parser to believe that error recovery is accomplished. The following statement in an action resets the parser to its normal mode:

```
yyerrok ;
```

The last example can be rewritten somewhat more usefully, as the following example shows:

```
input   :  error   '\n'
            {
               yyerrok;
               printf("Reenter last line: ");
            }
            input
        {
            $$ = $4;
        }
        ;
```

As previously mentioned, the token seen immediately after the `error` symbol is the input token at which the error was discovered. Sometimes this is inappropriate. For example, an error recovery action might take upon itself the job of finding the correct place to resume input. In this case, the previous look-ahead token must be cleared. The following statement in an action has this effect:

```
yyclearin ;
```

For example, suppose the action after `error` were to call some sophisticated resynchronization routine (supplied by the user) that attempted to advance the input to the beginning of the next valid statement. After this routine is called, the next token returned by `yylex` is presumably the first token in a legal statement. The old illegal token must be discarded and the error state reset. A rule similar to the following one could perform this:

```
stat    :  error
        {
           resynch();
           yyerrok  ;
           yyclearin;
        }
        ;
```

These mechanisms are admittedly crude but do allow for a simple, fairly effective recovery of the parser from many errors. Also, the user can get control to deal with the error actions required by other portions of the program.

The `yacc` environment

When the user enters a specification to `yacc`, the output is a file of C language programs called `y.tab.c`. The function produced by `yacc` is an integer-valued function called `yyparse`. When it is called, it in turn repeatedly calls `yylex`, the lexical analyzer supplied by the user (see "Lexical Analysis"), to obtain input tokens.

Eventually, if an error is detected, `yyparse` returns the value 1, and no error recovery is possible, or the lexical analyzer returns the end-marker token and the parser accepts. In this case, `yyparse` returns the value 0.

The user must provide a certain amount of environment for this parser to obtain a working program. For example, as with every C language program, a program called `main` must be defined that eventually calls `yyparse`. Also needed is a routine called `yyerror` that prints a message when a syntax error is detected. These two routines (`main` and `yyerror`) must be supplied in one form or another by the user.

To ease the initial effort of using `yacc`, a library is provided with default versions of `main` and `yyerror`. Use the `-ly` option of `ld` to incorporate these routines into your program. The following source code examples show the simplicity of these routines:

```
main()
{
   return ( yyparse() );
}
and
#include <stdio.h>

yyerror(s)
char *s;
{
   fprintf( stderr, "%s\n", s );
}
```

The argument to `yyerror` is a string containing an error message, usually the string `syntax error`. The average application wants to do better than this. Ordinarily, the program should keep track of the input line number and print it along with the message when a syntax error is detected.

The external integer variable `yychar` contains the look-ahead token number at the time the error was detected. This might be of some interest in giving better diagnostics.

Because the `main` program is probably supplied by the user (to read arguments, and so on), the `yacc` library is useful only in small projects or in the earliest stages of larger ones.

The external integer variable `yydebug` is normally set to 0. If it is set to a nonzero value, the parser sends as output a verbose description of its actions, including a discussion of the input symbols read and what the parser actions are. Depending on the operating environment, it might be possible to set `yydebug` by using a debugging system.

Input style

It is difficult to provide rules with substantial actions and still have a readable specification file. The following suggestions are a few style hints:

- Use all uppercase letters for token names and all lowercase letters for nonterminal names.
- Put grammar rules and actions on separate lines. This allows either to be changed without an automatic need to change the other.
- Put all rules with the same left side together. Put the left side in only once and let all following rules begin with a vertical bar.
- Put a semicolon only after the last rule with a given left side and put the semicolon on a separate line. This allows new rules to be easily added.
- Indent rule bodies by two tab stops and action bodies by three tab stops.

The example in "Example: A Desk Calculator" is written following this style (where space permits). You must make up your own mind about these stylistic questions. The central problem, however, is to make the rules visible through the morass of action code.

Left recursion

The algorithm used by the `yacc` parser encourages so-called *left recursive* grammar rules. Rules of the following form match this algorithm:

name : *name* *rest-of-rule* ;

Rules such as the following two frequently arise when writing specifications of sequences and lists. In each of these cases, the first rule is reduced for the first item only; the second rule is reduced for the second and all succeeding items:

```
list    :   item
            |   list ',' item
            ;

seq     :   item
            |   seq  item
            ;
```

With right recursive rules, such as the following examples, the parser is a bit bigger and the items are seen and reduced from right to left:

```
seq     :   item
        |   item   seq
        ;
```

More seriously, an internal stack in the parser is in danger of overflowing if a very long sequence is read. The user should use left recursion wherever reasonable.

It is worth considering whether a sequence with zero elements has any meaning; if so, consider writing the sequence specification as in the following, using an empty rule:

```
seq     :   /* empty */
        |   seq   item
;
```

Once again, the first rule is always reduced exactly once before the first item is read; the second rule is reduced once for each item read. Permitting empty sequences often leads to increased generality. However, conflicts might arise if `yacc` is asked to decide which empty sequence it has seen when it hasn't seen enough to know.

Lexical considerations

Some lexical decisions depend on context. For example, the lexical analyzer might want to delete blanks normally but not within quoted strings, or names might be entered into a symbol table in declarations but not in expressions.

One way of handling this situation is to create a global flag that is examined by the lexical analyzer and set by actions. The following example specifies a program that consists of zero or more declarations followed by zero or more statements. The flag `dflag` is 0 when reading statements and 1 when reading declarations, except for the first token in the first statement. This token must be seen by the parser before it can tell that the declaration section ended and the statements began. In many cases, this single token exception does not affect the lexical scan.

```
%{
   int dflag;
%}
```

... other declarations ...

```
%%

prog  :   decls stats
        ;

decls :   /* empty */
        {
            dflag = 1;
        }
        |   decls  declaration
        ;

stats :   /* empty */
        {
            dflag = 0;
        }
        |    stats  statement
        ;
```

... other rules ...

This kind of "back door" approach can be elaborated to an unpleasant degree. Nevertheless, it represents a way of doing some things that are difficult, if not impossible, to do otherwise.

Reserved words

Some programming languages permit you to use words (like `if`) that are normally reserved as label or variable names, provided that such use does not conflict with the legal use of these names in the programming language. This is extremely hard to do in the framework of `yacc`. It is difficult to pass information to the lexical analyzer telling it "this instance of `if` is a keyword and that instance is a variable." The user can try it using the mechanism described in the last section, but it is difficult. A number of ways of making this easier are being studied. For the time being, it is better that the keywords be reserved—that is, forbidden for use as variable names.

Simulating error and accept in actions

The parsing actions of error and accept can be simulated in an action by use of the macros `YYACCEPT` and `YYERROR`. The `YYACCEPT` macro causes `yyparse` to return the value 0. `YYERROR` causes the parser to behave as if the current input symbol had been a syntax error. The function `yyerror` is called, and error recovery takes place.

These mechanisms can be used to simulate parsers with multiple end-markers or context-sensitive syntax checking.

Accessing values in enclosing rules

An action can refer to values returned by actions to the left of the current rule. The mechanism is the same as with ordinary actions, a dollar sign followed by a digit.

```
sent    :   adj  noun  verb  adj  noun
        {
            look at the sentence ...
        }
        ;
```

```
adj     :   THE
        {
            $$ = THE;
        }
            | YOUNG
        {
            $$ = YOUNG;
        }
            ...
        ;
noun    :   DOG
        {
            $$ = DOG;
        }
          |   CRONE
        {
        if( $0 == YOUNG )
        {
            printf( "what?\n" );
        }
        $$ = CRONE;
       }
       ;
         ...
```

In this case, the digit can be 0 or negative.

In the action following the word `CRONE`, a check is made that the preceding token shifted was not `YOUNG`. Obviously, this is only possible when a great deal is known about what might precede the symbol `noun` in the input.

There also is a distinctly unstructured flavor about this. Nevertheless, at times, this mechanism prevents a great deal of trouble, especially when a few combinations are to be excluded from an otherwise regular structure.

Arbitrary value types

By default, the values returned by actions and the lexical analyzer are integers. The `yacc` program also can support values of other types, including structures. The `yacc` program keeps track of the types and inserts appropriate union member names so that the resulting parser is strictly type checked.

The `yacc` value stack is declared to be a union of the various types of values desired. The user declares the union and associates union member names to each token and nonterminal symbol having a value. When the value is referenced through a `$$` or `$n` construction, `yacc` automatically inserts the appropriate union name so that no unwanted conversions take place. This makes type-checking commands such as `lint` much quieter.

Three mechanisms are used to provide for this typing:

- First, there is a way of defining the union. This must be done by the user because other programs, notably the lexical analyzer, must know about the union member names.
- Second, there is a way of associating a union member name with tokens and nonterminal symbols.
- Third, there is a mechanism for describing the type of those few values where `yacc` cannot easily determine the type.

To declare the union, the user includes the following statement in the declaration section:

```
%union
{
    body of union
}
```

This declares the `yacc` value stack and the external variables `yylval` and `yyval` to have type equal to this union. If `yacc` was invoked with the `-d` option, the union declaration is copied onto the `y.tab.h` file. Alternatively, the union can be declared in a header file, and a `typedef` used to define the variable `YYSTYPE` to represent this union. Thus, the header file might have said the following, instead:

```
typedef union
{
    body of union
}
YYSTYPE;
```

The header file must be included in the declarations section by use of `%{` and `%}`. Once `YYSTYPE` is defined, the union member names must be associated with the various terminal and nonterminal names. The following construction is used to indicate a union member name:

<name>

If this follows one of the keywords `%token`, `%left`, `%right`, or `%nonassoc`, the union member name is associated with the tokens listed. For example, the following causes any reference to values returned by these two tokens to be tagged with the union member name `optype`:

```
%left  <optype>  '+'  '-'
```

Another keyword, `%type`, is used to associate union member names with nonterminals. For example, the following line can be used to associate the union member `nodetype` with the nonterminal symbols *expr* and `stat`.

`%type  <nodetype>  expr  stat`

There remain a couple of cases where these mechanisms are insufficient. If there is an action within a rule, the value returned by this action has no a priori type. Similarly, reference to left context values (such as `$0`) leaves `yacc` with no easy way of knowing the type. In this case, a type can be imposed on the reference by inserting a union member name between "<" and ">" immediately after the first `$`, as in the following example.

```
rule    :   aaa
        {
            $<intval>$ = 3;
        }
            bbb
        {
            fun( $<intval>2, $<other>0 );
        }
        ;
```

This syntax has little to recommend it, but the situation arises rarely. A sample specification is given in "Example: An Advanced Grammar." The facilities in this subsection are not triggered until they are used. In particular, the use of `%type` turns on these mechanisms. When they are used, there is a fairly strict level of checking. For example, use of `$n` or `$$` to refer to something with no defined type is diagnosed. If these facilities are not triggered, the `yacc` value stack is used to hold `int` values, as was true historically.

Example: A desk calculator

This section contains an example that gives the complete `yacc` applications for a small desk calculator. The calculator has 26 registers labeled `a` through `z` and accepts arithmetic expressions made up of the operators shown in Table 3-2.

If an expression at the top level is an assignment, the value is printed. Otherwise, the expression is printed. As in the C language, an integer that begins with 0 (zero) is assumed to be octal. Otherwise, it is assumed to be decimal.

As an example of a `yacc` specification, the desk calculator does a reasonable job of showing how precedence and ambiguities are used and demonstrates simple recovery. The major oversimplifications are that the lexical analyzer is much simpler than what is necessary for most applications, and the output is produced immediately line by line.

Table 3-2 Arithmetic operators

Symbol	Meaning
+	Addition
-	Subtraction
*	Multiplication
/	Division
%	Modulus (Remainder)
&	Binary AND
|	Binary OR
=	Assignment

Note the way that decimal and octal integers are read in by grammar rules. This job is probably better done by the lexical analyzer.

```
%{
#include <stdio.h>
#include <ctype.h>

int regs[26];
int base;

%}
%start list
%token DIGIT LETTER
%left '|'
%left '&'
%left '+'  '-'
%left '*'  '/'  '%'
%left UMINUS     /* precedence for unary minus */
%%               /* beginning of rule section */
```

(continued)➡

```
list    :       /* empty */
        |       list stat ' \n'
        |       list error ' \n'
        {
                yyerror;
        }
        ;
stat    :       expr
        {
                printf( "%d\n", $1 );
        }
        |       LETTER '='  expr
        {
                regs[$1] = $3
        }
        ;
expr    :       '(' expr ')'
        {
                $$ = $2;
        }
        |       expr '+' expr
        {
                $$ = $1 + $3
        }
        |       expr '-' expr
        {
                $$ = $1 - $3
        }
        |       expr '*' expr
        {
                $$ = $1 * $3;
        }
        |       expr '/' expr
```

```
        {
                $$ = $1/$3;
        }
        |       exp '%' expr
        {
                $$ = $1 % $3
        }
        |       expr '&' expr
        {
                $$ = $1 & $3;
        }
        |       expr '|' expr
        {
                $$ = $1 | $3
        }
        |       '-' expr  %prec UMINUS
        {
                $$ = - $2;
        }
        |       LETTER
        {
                $$ = reg[$1];
        }
        |       number
        ;
number :     DIGIT
        {
                $$ = $1; base = ($1==0) ? 8 : 10;
        }
        |       number DIGIT
        {
                $$ = base * $1 + $2
        }
       ;
```

(continued)➡

```
%% /* start of program */
/*
 * lexical analysis routine
 * return LETTER for lowercase letter
 *     (i.e., yylval = 0 through 25)
 * returns DIGIT for digit
 *     (i.e., yylval = 0 through 9)
 * all other characters are returned immediately
 *
 */
yylex( )
{
  int c;
    while (c=getchar( )) == '  ') /* skip blanks */
        ;
        if( islower( c ))
        {
            yylval = c - 'a' ;
                return( LETTER );
        }
        if( isdigit( c ))
        {
            yylval = c - '0' ;
                return( DIGIT );
        }
  return( c );
}
```

Example: `yacc` input syntax

This section contains a description of the `yacc` input syntax as a `yacc` specification. Context dependencies, and so forth, are not considered. Ironically, the `yacc` input specification language is most naturally specified as an LR(2) grammar. The sticky part comes when an identifier is seen in a rule immediately following an action. If this identifier is followed by a colon, it is the start of the next rule; otherwise, it is a continuation of the current rule, which just happens to have an action embedded in it.

As implemented, the lexical analyzer looks ahead after seeing an identifier and decides whether the next token (skipping blanks, newlines, comments, and so on) is a colon. If so, it returns the token `C_IDENTIFIER`. Otherwise, it returns `IDENTIFIER`. Literals (quoted strings) are also returned as `IDENTIFIER` but never as part of `C_IDENTIFIER`.

```
/* grammar for the input to yacc */

/* basic entries */
/* includes identifiers and literals */
%token IDENTIFIER
/* identifier (but not literal) followed by a colon */
%token C_IDENTIFIER
%token NUMBER          /* [0-9]+ */

/* reserved words: */
/* %type -> TYPE,  %left -> LEFT,   etc. */
%token LEFT RIGHT NONASSOC TOKEN PREC TYPE START UNION
%token MARK            /* the %% mark */
%token LCURL           /* the %{ mark */
%token RCURL           /* the %} mark */

/* ASCII character literals stand for themselves */
%token spec

%%
```

(continued)➡

```
spec    :   defs MARK rules tail
        ;
tail    :   MARK
        {
            ...In this action, read the rest of the file...
        }
        |   /* empty: the second MARK is optional */
        ;
defs    :   /* empty */
        |   defs def
        ;
defs    :   START IDENTIFIER
        |   UNION
        {
            ...Copy union definition to output...
        }
        |   LCURL
        {
            ...Copy C code to output file...
            RCURL
        }
        |   ndefs rword tag nlist
        ;
rword   :   TOKEN
        |   LEFT
        |   RIGHT
        |   NONASSOC
        |   TYPE
        ;
tag     :   /* empty: union tag is optional */
        |   '<' IDENTIFIER '>'
        ;
```

```
nlist   :   nmno
        |   nlist nmno
        |   nlist ',' nmno
        ;
/* Note: literal illegal with %type */
nmno    :   IDENTIFIER
        |   IDENTIFIER NUMBER
        ;
/* rule section */
rule    :   C_IDENTIFIER rbody proc
        |   rule rule
        ;
rule    :   C_IDENTIFIER rbody prec
        |   '|' rbody prec
        ;
rbody   :   /* empty */
        |   rbody IDENTIFIER
        |   rbody act
        ;
act     :   '{'
        {
            ...Copy action, translate $$'s etc....
        }
            '}'
        ;
prec    :   /* empty */
        |   PREC IDENTIFIER
        |   PREC IDENTIFIER act
        |   prec';'
        ;
```

Example: An advanced grammar

This section gives an example of a grammar using some of the advanced features. It modifies the example from "Example: A Desk Calculator" to provide a desk calculator that does floating-point interval arithmetic.

The calculator understands floating-point constants, as well as the arithmetic operations `+`, `-`, `*`, `/`, unary `-`, and the letters `a` through `z`. The calculator also understands intervals written as is the following example, where `X` is less than or equal to `Y`:

```
(X,Y)
```

There are 26 interval valued variables `A` through `Z` that can also be used. The usage is similar to that in "Example: A Desk Calculator." That is, assignments return no value and print nothing, while expressions print the floating or interval value.

Intervals are represented by a structure consisting of the left and right endpoint values stored as doubles. This structure is given a type name, `INTERVAL`, by using `typedef`. The `yacc` value stack can also contain floating-point scalars and integers that are used to index into the arrays holding the variable values. The entire strategy depends strongly on being able to assign structures and unions in the C language. In fact, many of the actions call functions that return structures as well.

Note the use of `YYERROR` to handle error conditions: division by an interval containing 0 and an interval presented in the wrong order. The error-recovery mechanism of `yacc` is used to throw away the rest of the offending line. In addition to the mixing of types on the value stack, this grammar also demonstrates an interesting use of syntax to keep track of the type (for example, scalar or interval) of intermediate expressions. Scalars can be automatically promoted to an interval if the context demands an interval value. This causes a large number of conflicts when the grammar is run through `yacc`—18 shift/reduce and 26 reduce/reduce. The problem can be seen by looking at the following input lines:

```
2.5+(3.5-4.)
```

and

```
2.5 + ( 3.5,4 )
```

Notice that the `2.5` is to be used in an interval-value expression in the second example, but this fact is not known until the comma is read. By this time `2.5` is finished, and the parser cannot go back and change its mind.

More generally, it might be necessary to look ahead an arbitrary number of tokens to decide whether to convert a scalar to an interval. This problem is evaded by having two rules for each binary interval valued operator, one when the left operand is a scalar and one when the left operand is an interval. In the second case, the right operand must be an interval, so the conversion is applied automatically.

Despite this evasion, there are still many cases where the conversion might be applied or not, leading to the above conflicts. They are resolved by listing the rules that yield scalars first in the specification file. In this way, the conflict is resolved in the direction of keeping scalar-valued expressions scalar valued until they are forced to become intervals. This way of handling multiple types is very instructive, but not very general. If there were many kinds of expression types instead of just two, the number of rules needed would increase dramatically and the conflicts even more dramatically. Thus, while this example is instructive, it is better practice in a more normal programming language environment to keep the type information as part of the value and not as part of the grammar.

Finally, a word about the lexical analysis. The only unusual feature is the treatment of floating-point constants. The C language library routine `atof` is used to do the actual conversion from a character string to a double-precision value. If the lexical analyzer detects an error, it responds by returning a token that is illegal in the grammar, provoking a syntax error in the parser and thence error recovery.

```
%{

#include<stdio.h>
#include<ctype.h>

typedef struct interval
{
   double lo, hi;
}  INTERVAL;

INTERVAL vmul(), vdiv( );

double atof();
double dreg[26];
INTERVAL vreg[26];

%}
```

(continued)➡

```
%start line
%union
{
  int ival;
  double dval;
  INTERVAL vval;
}

%token <ival> DREG VREG /*indices into dreg, vreg */
%token <dval> CONST     /* floating point constant */

%type <dval> dexp       /* expression */
%type <vval> vexp       /* interval expression */

  /* precedence information about the operators */

%left   '+'  '-'
%left   '*'  '/'
%left   UMINUS   /* precedence for unary minus */

%%
lines   :  /* empty */
        |  lines line
        ;
line    :  dexp '\n'
        {
          printf( "%15.8f\n".$1 );
        }
        |  vexp '\n'
        {
          printf("(%15.8f,%15.8f)\n",$1.1o,$1.hi );
        }
        |  DREG '='  '\n'
        {
          dreg[$1] = $3;
        }
        |  VREG '='  vexp '\n'
```

```
                {
                  vreg[$1] = $3;
                }
                |  error '\n'
                {
                  yyerrork;
                }
                ;
dexp            :  CONST
                |  DREG
                {
                  $$ = dreg[$1]
                }
                |  dexp '+' dexp
                {
                  $$ = $1 + $3
                }
                |  dexp '-' dexp
                {
                  $$ = $1 - $3
                }
                |  dexp '*' dexp
                {
                  $$ = $1 * $3
                }
                |  dexp '/' dexp
                {
                  $$ = $1 / $3
                }
                | '-'  dexp    %prec UMINUS
                {
                  $$ =- $2
                }
```

(continued)➡

```
        |  '(' dexp ')'
        {
          $$ = $2 .
        }
        ;
vexpp   :  dexp
        {
          $$.hi = $$.lo = $1;
        }
        |  '(' dexp ',' dexp ')'
        {
          $$.lo = $2;
          $$.hi = $4;
          if( $$.lo > $$.hi )
          {
            printf( "interval out of order n" );
            YYERROR;
          }
        }
        |   VREG
        {
          $$ = vreg[$1]
        }
        |  vexp '+' vexp
        {
          $$.hi = $1.hi + $3.hi;
          $$.lo = $1.lo + $3.lo
        }
        |  dexp '+' vexp
        {
          $$.hi = $1 + $3.hi;
          $$.lo = $1 + $3.lo
        }
```

```
|  vexp '=' vexp
{
  $$.hi = $1.hi - $3.lo;
  $$.lo = $1.lo - $3.hi
}
|  dvep '-'  vdep
{
  $$.hi = $1 - $3.lo;
  $$.lo = $1 - $3.hi
}
|  vexp '*' vexp
{
  $$ = vmul( $1.lo,$.hi,$3 )
}
|  dexp '*' vexp
{
  $$ = vmul( $1, $1, $3 )
}
|  vexp '/' vexp
{
  if( dcheck( $3 ) ) YYERROR;
  $$ = vdiv( $1.lo, $1.hi, $3 )
}
|  dexp '/' vexp
{
  if( dcheck( $3 ) ) YYERROR;
  $$ = vdiv( $1.lo, $1.hi, $3 )
}
|  '-' vexp    %prec UMINUS
{
  $$.hi = -$2.lo;$$.lo =-$2.hi
}
|  '(' vexp ')'
```

(continued)➡

```
        }
          $$ = $2
        }
        ;
%%
/* buffer size for floating point number */
# define BSZ 50
/*
 *lexical analysis
 */
yylex( )
{
    register c;
    while ((c=getchar()) == ' ' ) /* skip blanks */ ;
    if(isupper(c))
    {
        yylvalval = c - 'A'
        return(VREG);
    }
    if(islower(c))
    {
        yylvalval = c - 'a' ,
        return(DREG);
    }
/*
 * gobble up digits, points, exponents
 */
    if(isdigit(c) || c == ' .' )
    {
        char buf[BSZ+1], *cp = buf;
        int dot = 0, exp = 0;
```

```
            for(;(cp - buf) < BSZ ; ++cp,c=getchar())
            {
                *cp = c;
                if(isdigit(c))
                  continue;
                if(c == '.' )
            {
                if(dot++ || exp)
                  /* causes syntax error */
                   return( '.' );
                continue;
            }
            if(c == 'e' )
            {
              if( exp++ )
                 /* causes syntax error */
                 return( 'e' );
              continue;
            }
             break;                /* end of number */
        }
        *cp = '\0' ;
        if((cp - buff) >= BSZ)
            printf( "constant too long truncated\n");
        else
            /* push back last char read */
            ungetc(c, stdin);
        yylval.dval = atof(buf);
        return(CONST);
      }
      return(c);
    }
```

(continued)➡

```
/*
 * returns the smallest interval
 * between a, b, c and d
 */

INTERVAL hilo( a, b, c, d )
double a, b, c, d;
{
    INTERVAL v;
    if( a>b )
    {
        v.hi = a;
        v.lo = b;
    }
    else
    {
        v.hi = b;
        v.lo = a;
    }
    if( c>d )
    {
        if( c>v.hi )
           v.hi = c;
        if( d<v.lo )
           v.lo = d;
    }
    else
    {
        if( d>v.hi )
           v.hi = d;
        if( c<v.lo )
           v.lo = c;
    }
```

```
    return( v );
}
INTERVAL vmul( a, b, v )
double a, b;
INTERVAL v;
{
    return( hilo( a*v.hi, a*v.lo, b*v.hi, b*v.lo ) );
}
dcheck( v )
INTERVAL v;
{
    if( v.hi >=0.&& v.lo <=0. )
    {
       printf( "divisor internal contains 0.\n" );
       return( 1);
    }
    return( 0 );
}
INTERVAL vdiv( a, b, v )
double a, b;
INTERVAL v;
{
    return( hilo( a/v.hi, a/v.lo, b/v.hi, b/v.lo ) );
}
```

Backward compatibility

This section mentions synonyms and features that are supported for historical continuity but, for various reasons, are not encouraged.

- Literals can also be delimited by double quotes.
- Literals can be more than one character long. If all the characters are alphabetic, numeric, or _, the type number of the literal is defined just as if the literal did not have the quotes around it. Otherwise, it is difficult to find the value for such a literal.

  The use of multicharacter literals is likely to mislead those unfamiliar with `yacc`, because it suggests that `yacc` is doing a job that actually must be done by the lexical analyzer.
- Most places where (`%`) is legal, the backslash (`\`) can be used. In particular, `\\` is the same as `%%`, `\left` the same as `%left`, and so on.
- There are a number of other synonyms:

`%<`	is the same as	`%left`
`%>`	is the same as	`%right`
`%binary`	is the same as	`%nonassoc`
`%2`	is the same as	`%nonassoc`
`%0`	is the same as	`%token`
`%term`	is the same as	`%token`
`%=`	is the same as	`%prec`

- Actions can also have the form

  ```
  ={ ... }
  ```

  and the braces can be dropped if the action is a single C language statement.
- C language code between `%{` and `%}` used to be permitted at the head of the rules section as well as in the declaration section.

4 m4: A Macro Processor

Invoking m4 / 4-3
Defining macros / 4-3
Arithmetic built-ins / 4-9
I/O manipulation / 4-10
String manipulation / 4-14
Printing / 4-16
Executing system commands / 4-16
Interactive use of m4 / 4-17
Recursive definitions / 4-17
Built-in macro summary / 4-19

The m4 macro processor is a general-purpose macro-processing utility. It can also be considered to be an interpreter for the m4 language. The #define statement in the C language is an example of the basic facility provided by any macro processor: the replacement of some text by some (other) text. For several reasons, m4 is a more powerful macro processor than the standard C preprocessor, cpp.

The basic operation of `m4` is to read every alphanumeric token (string of letters and digits) in the input and to determine whether the token is the name of a macro. The name of a macro is replaced by its defining text and the resulting string is pushed back onto the input to be rescanned.

In addition to the straightforward replacement of one string of text by another, the `m4` macro processor also provides the following features:

- arguments to macros
- arithmetic capabilities
- file manipulation
- conditional macro expansion
- string and substring functions
- recursive definitions

When a macro is called with arguments, the arguments are collected and substituted into the right places in the defining text before the defining text is rescanned.

The `m4` macro processor accepts user-defined macros as well as its "built-in" macros. Both types of macros work exactly the same way, except that some of the built-in macros have side effects on the state of the process.

Invoking `m4`

To run `m4`, give the command

`m4` *files*

Each argument file is processed in order. If there are no arguments, or if an argument is `-`, the standard input is read at that point.

The processed text is written on the standard output. The output can be redirected for subsequent processing, as follows:

`m4` *files* `>` *outputfile*

Defining macros

The `m4` macro allows you to define macros, remove their definition, have conditional definitions, specify arguments in a definition, and many other tasks. This section outlines important aspects of macro definition.

`define`

The primary built-in function of `m4` is `define`. This function is used to define new macros. The general form is

`define(`*name*`,` *replacement*`)`

All subsequent occurrences of *name* are replaced by *replacement*. The *name* must be alphanumeric and must begin with a letter (the underscore (_) counts as a letter). The *replacement* is any text that contains balanced parentheses. An escaped RETURN or an embedded newline character allows a multiline *replacement* to be specified.

The following is a typical example of the use of `define`, in which `N` is defined to be the string `100` and is then used in a later `if` statement:

```
define(N, 100)
if (i > N) echo "number too large"
```

The left parenthesis must immediately follow the word `define` to signal that `define` has arguments. If a user-defined macro or built-in name is not followed immediately by this character, the macro call is assumed to have no arguments.

Macro calls have the following general form:

name`(`*arg1*`,` *arg2*`, ...,` *argn*`)`

A macro name is recognized as such only if it appears surrounded by nonalphanumerics. In the following example, the variable `NNN` is absolutely unrelated to the defined macro `N`, even though the variable contains a lot of `N`'s:

```
define(N, 100)
if (NNN > 100) echo "number too large"
```

Macros can be defined in terms of other macros. For example, the following defines both `M` and `N` to be `100`. If `N` is redefined and subsequently changes, `M` retains the value of `100`, not `N`.

```
define(N, 100)
define(M, N)
```

The `m4` macro processor expands macro names into their defining text as soon as possible. The string `N` is immediately replaced by `100`. The string `M` is then defined to be 100. The overall result is the same as using the following input in the first place:

```
define(M, 100)
```

The order of the definitions can be interchanged as follows:

```
define(M, N)
define(N, 100)
```

Now `M` is defined to be the string `N`, so when the value of `M` is requested later, the result is the value of `N` at that time (because the `M` is replaced by `N`, which is replaced by `100`).

Quoting

The more general solution to the problem of making sure the correct strings get substituted is to delay the expansion of the arguments of `define` by quoting them. The quoting characters initially recognized by `m4` are the left and right single quotes, (`` ` `` and `'`). Any text surrounded by left and right single quotes is not expanded immediately but has the quotes stripped off. The value of a quoted string is the string stripped of the quotes. If the input is

```
define(N, 100)
define(M, `N')
```

the quotes around the `N` are stripped off as the argument is being collected. The result of using quotes is to define `M` as the string `N`, not as `100`.

The general rule is that `m4` always strips off one level of single quotes whenever it evaluates something. *This is true even outside macros.*

If the word `define` itself is to appear in the output, the word must be quoted in the input as follows:

```
`define' = 1;
```

Another example of using quotes is to redefine a macro. To redefine `N`, the evaluation must be delayed by quoting:

```
define(N, 100)
define(`N', 200)
```

In `m4`, it is often wise to quote the first argument of a macro. The following example, for instance, does not redefine `N`:

```
define(N, 100)
define(N, 200)
```

The `N` in the second definition is replaced by `100`. The result is equivalent to the following statement:

```
define(100, 200)
```

This statement is ignored by `m4`, however, because only names that begin with an alphanumeric character can be defined.

changequote

If left and right single quotes are not convenient for some reason, the quote characters can be changed with the following built-in macro:

```
changequote([, ])
```

The built-in `changequote` makes the new quote characters the left and right brackets. The original characters can be restored by using `changequote` without arguments, as follows:

```
changequote
```

undefine

The `undefine` macro removes the definition of some macro or built-in as follows:

```
undefine(`N')
```

The macro removes the definition of `N`. Built-ins can be removed with `undefine`, as follows:

```
undefine(`define')
```

Once removed, the definition cannot be reused.

ifdef

The built-in `ifdef` provides a way to determine whether a macro is currently defined.

Depending on the system, a definition appropriate for the particular machine can be made as follows:

```
ifdef(`pdp11', `define(wordsize,16)')
ifdef(`u3b', `define(wordsize,32)')
```

Remember to use the quotes.

The `ifdef` macro actually permits three arguments. If the first argument is defined, the value of `ifdef` is the second argument. If the first argument is not defined, the value of `ifdef` is the third argument. If there is no third argument, the value of

`ifdef` is null. If the name is undefined, the value of `ifdef` is then the third argument, as in

```
ifdef('unix', on UNIX, not on UNIX)
```

Arguments

User-defined macros can also have arguments, so different invocations can have different results. Within the replacement text for a macro (the second argument of its `define`), any occurrence of `$`*n* is replaced by the *n*th argument when the macro is actually used. Thus, the following macro, `bump`, generates code to increment its argument by 1:

```
define(bump, $1 = $1 + 1)
```

The statement

```
bump(x)
```

is equivalent to

```
x = x + 1
```

A macro can have as many arguments as needed, but only the first nine are accessible (`$1` through `$9`) (see "Built-In Macro Summary" under `shift` for more information). The macro name is `$0`, although that is less commonly used. Arguments that are not supplied are replaced by null strings, so a macro can be defined that simply concatenates its arguments like this:

```
define(cat, $1$2$3$4$5$6$7$8$9)
```

Thus,

```
cat(x, y, z)
```

is equivalent to

```
xyz
```

Arguments `$4` through `$9` are null, because no corresponding arguments are provided. Leading unquoted blanks, tabs, or newlines that occur during argument collection are discarded. All other white space is retained. Thus,

```
define(a, b c)
```

defines `a` to be `b c`.

Arguments are separated by commas; however, when commas occur within parentheses, the argument is neither terminated nor separated. For example,

```
define(a, (b,c))
```

has only two arguments. The first argument is `a`. The second is literally `(b,c)`. A bare comma or parenthesis can be inserted by quoting it.

Three other constructions are useful in macro definitions:

```
$#
$*
$@
```

During macro replacement, the construction `$#` is replaced by the number of arguments. The `$*` construction is replaced by a list of the arguments separated by commas. The construction `$@` is like `$*` except that each argument is quoted (using the current quotes). See the section "Recursive Definitions" for examples of the first two constructions.

ifelse

Arbitrary conditional testing is performed through the built-in macro `ifelse`. In the simplest form,

`ifelse(`*a*`, `*b*`, `*c*`, `*d*`)`

compares the two strings *a* and *b*. If *a* and *b* are identical, `ifelse` returns the string *c*. Otherwise, string *d* is returned. Thus, a macro called `compare` can be defined to compare two strings and return `yes` or `no` if they are the same or different, as follows:

```
define(compare, `ifelse($1, $2, yes, no)')
```

Note the quotes, which prevent evaluation of `ifelse` occurring too early. If the fourth argument is missing, it is treated as empty. Thus,

`ifelse(`*a*`, `*b*`, `*c*`)`

is *c* if *a* matches *b*, and null otherwise.

`ifelse` can actually have any number of arguments and provides a limited form of multiway decision capability. In the input

`ifelse(`*a*`, `*b*`, `*c*`, `*d*`, `*e*`, `*f*`, `*g*`)`

if the string *a* is the same as the string *b*, the result is *c*. Otherwise, if *d* is the same as *e*, the result is *f*. Otherwise, the result is *g*. If the final argument is omitted and the specified strings don't match, the result is null.

Arithmetic built-ins

The `m4` program provides three built-in functions for doing arithmetic on integers (only):

```
incr
decr
eval
```

The simplest are `incr`, which increments its numeric argument by 1, and `decr`, which decrements by 1. Thus, to handle the common programming situation where a variable is to be defined as "one more than `N`," use the following form:

```
define(N, 100)
define(N1, `incr(N)')
```

Then `N1` is defined as one more than the current value of `N`.

The more general mechanism for arithmetic is a built-in function called `eval`, which is capable of arbitrary arithmetic on integers. The operators in decreasing order of precedence are shown in Table 4-1.

Parentheses can be used to group operations where needed. All the operands of an expression given to `eval` must ultimately be numeric. The numeric value of a true relation (like `1>0`) is 1 and false is 0. The precision in `eval` is 32 bits under the A/UX operating system.

Table 4-1 Arithmetic operators

Symbol	Meaning
`+ -`	Unary plus and minus
`**^`	Exponentiation
`* / %`	Multiplication and division
`+-`	Binary plus and minus
`== != < <= > >=`	Relational operators
`!`	Logical negation (NOT)
`& &&`	Logical multiplication (AND)
`| ||`	Logical addition (OR)

As a simple example, define `M` to be `2==N+1` using `eval` as follows:

```
define(N, 3)
define(M, `eval(2==N+1)')
```

First `N` is defined as 3; then `M` is defined as 0, since 2 is not equal to `N+1`. If `M` were defined as

```
define(M, `eval(2==N-1)')
```

then its defined value would be 1, because the result of the comparison would be true.

The defining text for a macro should be quoted unless the text is very simple. Quoting the defining text usually gives the desired result and is a good habit to get into.

I/O manipulation

The `m4` utility provides numerous functions to handle input and output. These routines are detailed in this section.

`include` and `sinclude`

A new file can be included in the input at any time by the built-in function `include`. For example,

`include` (*filename*)

inserts the contents of *filename* in place of the `include` command. The contents of the file are often a set of definitions. The value of `include` (the replacement text of `include`) is the contents of the file. If needed, the contents can be captured in definitions, and so on. A fatal error occurs if the file named by *filename* cannot be accessed. To get some control over this situation, you can use the alternate form, `sinclude`, or quote the filename. The built-in `sinclude` (silent include) says nothing and continues if the file named cannot be accessed.

`divert`, `undivert`, and `divnum`

The output of `m4` can be diverted to temporary files during processing, and the collected material can be generated upon command. The `m4` program maintains nine of these diversions, numbered 1 through 9. If the built-in macro

`divert(`*n*`)`

is used, all subsequent output is put onto the end of a temporary file referred to as *n*. Diverting to this file is stopped by the `divert` or `divert(0)` command, which resumes the normal output process.

Diverted text is normally produced all at once at the end of processing with the diversions produced in ascending numerical order. Diversions can be brought back at any time by appending the new diversion to the current diversion. Output diverted to a stream other than 0 through 9 is discarded. The following code, for example, throws away excess newlines:

```
divert(-1)
define(N, 100)
define(M, 200)
define(L, 300)
divert
```

◆ **Note** The newline character at the end of each `define` is passed to the output, as described in the following section. ◆

The built-in macro `undivert`, with no arguments, brings back all diversions in numerical order. With arguments, `undivert` brings back the selected diversions in the order specified by the argument. `undivert` discards the diverted text. You can also discard text by using a diversion number that is not between 0 and 9, inclusive.

The value of `undivert` is *not* the diverted text, but rather the number of the diversion to bring back into the text. Furthermore, the diverted material is not rescanned for macros.

As an example of the interaction between `divert`, `undivert`, and current diversion, consider the following code:

```
this is current diversion
divert(1)
this is diversion 1
divert(2)
this is diversion 2
divert(3)
this is diversion 3
divert
this is current diversion again
undivert
once again, current diversion
```

In the above trivial code there are three diversions between the two lines of current diversion code. The use of `divert` at the end of diversion 3 is needed to inform `m4` that what follows is not part of diversion 3. `undivert` with no arguments inserts at the current position all previous diversions, with no rescanning of any macros that might be there. The output of the above code is

```
this is current diversion

this is current diversion again

this is diversion 1

this is diversion 2

this is diversion 3

once again, current diversion
```

Note that the diverted text is not brought back again at the end of the output by the normal process; the diverted text is discarded by the use of `undivert`. Another example can make this clearer:

```
this is main diversion
divert(1)
this is diversion 1
divert(2)
this is diversion 2
divert(3)
this is diversion 3
divert
this is main diversion again
undivert(3)
once again, main diversion
undivert(2)
```

The ouput for the above is

```
this is main diversion

this is main diversion again

this is diversion 3

once again, main diversion

this is diversion 2

this is diversion 1
```

As you can see, only diversion 1 is brought back by the normal process, because only diversion 1 is not undiverted and, therefore, discarded. Note also that you can change the order of appearance of the diverted versions.

The built-in macro `divnum` returns the number of the currently active diversion. The current output stream is 0 during normal processing.

dnl

There is a built-in macro called `dnl` that deletes all characters that follow it, up to and including the next newline. The `dnl` macro is useful mainly for throwing away empty lines that otherwise tend to clutter up `m4` output. Using input

```
define(N, 100)
define(M, 200)
define(L, 300)
```

results in a newline at the end of each line that is not part of the definition. The newline is copied into the output so that each `define` statement is followed by a blank line. If the built-in macro `dnl` is added to each of these lines, the newlines disappear.

```
define(N, 100)dnl
define(M, 200)dnl
define(L, 300)dnl
```

String manipulation

The `m4` utility provides numerous functions to handle string manipulation. These routines are detailed in this section.

len

The built-in macro `len` returns the length of the string (number of characters) that makes up its argument. Thus,

```
len(abcdef)
```

is 6, and

```
len((a,b))
```

is 5 (the parentheses and comma are counted along with `a` and `b`).

substr

The built-in macro `substr` can be used to produce substrings of strings. The input

`substr(`*s*`, `*i*`, `*n*`)`

returns the substring of *s* that starts at the *i*th position (origin 0) and is *n* characters long. If *n* is omitted, the rest of the string is returned. For example,

```
substr('now is the time',1)
```

returns the following string:

```
ow is the time.
```

If *i* or *n* is out of range, various actions occur.

index and translit

The built-in macro `index` returns the index (position) in one string where the first character of another given string occurs, or –1 if it does not occur. It is written as

`index(`*s1*`, `*s2*`)`

where *s1* is the string to be searched and *s2* is the string to be searched for. As with `substr`, the origin for strings is 0.

The built-in macro `translit` performs character transliteration and has the general form

`translit(`*s*`, `*f*`, `*t*`)`

which modifies *s* by replacing any character found in *f* by the corresponding character of *t*. Using

```
translit(s, aeiou, 12345)
```

replaces the vowels by the corresponding digits. If *t* is shorter than *f*, characters that do not have an entry in *t* are deleted. As a limiting case, if *t* is not present at all, characters from *f* are deleted from *s*. So,

```
translit(s, aeiou)
```

deletes vowels from *s*.

Printing

This section details the `m4` routines for printing.

`errprint`

The built-in macro `errprint` writes its arguments out on the standard error file. An example is

```
errprint('fatal error')
```

`dumpdef`

The built-in macro `dumpdef` is a debugging aid that dumps the current names and definitions of items named as arguments. If no arguments are given, then all current names and definitions are printed. Remember to quote the names.

Executing system commands

This section describes the `m4` routines that execute system commands.

`syscmd` and `maketemp`

Any program in the local operating system can be run by using the built-in macro `syscmd`. For example,

```
syscmd(date)
```

on the A/UX system runs the `date` command. Normally, `syscmd` is used to create a file for a subsequent `include`.

To facilitate making unique filenames, the built-in macro `maketemp` is provided with specifications identical to the system function `mktemp`. The `maketemp` macro fills in a string of `XXXXX` in the argument with the process ID of the current process.

Interactive use of `m4`

The input to `m4` can come from a file, the standard input, or both. Thus, it is possible to use `m4` interactively, by telling it to take its input from the standard input. There are several ways to do this. The simplest is to invoke `m4` as follows:

```
m4
```

At this point, `m4` reads from the standard input.

If you have an existing set of `m4` commands stored in a file, you can instruct `m4` to process those commands first by invoking it as

`m4` *file* `-`

The minus sign is required here to instruct `m4` to read *file* and then the standard input. Alternatively, if you invoke `m4` using just the `m4` command with no arguments, you can tell `m4` to fetch the set of commands from *file* by typing the following line:

`include(`*file*`)`

The effect is the same in both cases.

Recursive definitions

Since `m4` rescans any text that arises from the replacement of a macro by its defining text, it is possible to construct recursive macro definitions. That is, it is perfectly legal to define a macro in terms of itself. As with any well-constructed recursive definition, however, you must take care that the definition has a well-defined stopping point. Generally, this is easy to do with the `ifelse` command.

For instance, suppose that you need a macro that returns its last argument and discards the rest. You might write the following definition:

```
define(last,
`ifelse($#,1,$1, `last(shift($*))')')
```

When there are multiple arguments, `last` drops the first argument and then calls itself to look for the last argument in the remaining argument list. This definition is well behaved because when there is only one argument, it alone is returned.

A more interesting example is the following definition of the factorial function:

```
define(fact,
`ifelse($1,1,1, `eval($1*fact(decr($1)))')')
```

If you give `m4` the following input,

```
The factorial of 1 is fact(1).
The factorial of 2 is fact(2).
The factorial of 3 is fact(3).
The factorial of 4 is fact(4).
The factorial of 5 is fact(5).
The factorial of 6 is fact(6).
The factorial of 7 is fact(7).
The factorial of 8 is fact(8).
```

you get the following output:

```
The factorial of 1 is 1.
The factorial of 2 is 2.
The factorial of 3 is 6.
The factorial of 4 is 24.
The factorial of 5 is 120.
The factorial of 6 is 720.
The factorial of 7 is 5040.
The factorial of 8 is 40320.
```

Finally, you might want to define a recursive macro with two arguments. The standard power function serves nicely:

```
define(pow,
`ifelse($2,1,$1, `eval($1*pow($1,decr($2)))')')
```

If you then give `m4` the following input,

```
3 to power 1 is pow(3,1).
3 to power 2 is pow(3,2).
3 to power 3 is pow(3,3).
3 to power 4 is pow(3,4).
3 to power 5 is pow(3,5).
3 to power 6 is pow(3,6).
```

```
3 to power 7 is pow(3,7).
3 to power 8 is pow(3,8).
```

you get

```
3 to power 1 is 3.
3 to power 2 is 9.
3 to power 3 is 27.
3 to power 4 is 81.
3 to power 5 is 243.
3 to power 6 is 729.
3 to power 7 is 2187.
3 to power 8 is 6561.
```

Built-in macro summary

The following items are `m4` built-in macros:

`changecom`	Changes left and right comment markers from the default `#` and newline. With no arguments, the comment mechanism is disabled. Comment markers can be up to five characters long.
`changequote`	Changes quoting symbols to the first and second arguments. The symbols can be up to five characters long. With no arguments, this macro restores the original quote characters.
`decr`	Returns the value of its argument decremented by 1.
`define`	Defines new macros.
`defn`	Returns the quoted definition of its arguments.
`divert`	Diverts output to one of ten diversions (named 0 through 9).
`divnum`	Returns the number of the currently active diversion.
`dnl`	Reads and discards characters up to and including the next newline.
`dumpdef`	Dumps the current names and definitions of items named as arguments. With no arguments, definitions of all current macros are dumped.
`errprint`	Prints its arguments on the standard error file.
`eval`	Performs arbitrary arithmetic on integers.

`ifdef`	Determines whether a macro is currently defined.
`ifelse`	Performs arbitrary conditional testing.
`include`	Returns the contents of the file named in the argument. A fatal error occurs if the file named cannot be accessed.
`incr`	Returns the value of its argument incremented by 1.
`index`	Returns the position where the second argument begins in the first argument.
`len`	Returns the number of characters that make up its argument.
`m4exit`	Causes immediate exit from `m4`.
`m4wrap`	Pushes the exit code back at final end-of-file (EOF).
`maketemp`	Facilitates making unique filenames.
`popdef`	Removes the current definition of its arguments, exposing any previous definitions.
`pushdef`	Defines new macros but saves any previous definition.
`shift`	Returns all arguments except the first argument.
`sinclude`	Returns the contents of the file named in the arguments. The macro remains silent and continues if the file is inaccessible.
`substr`	Produces substrings of strings.
`syscmd`	Executes the A/UX system command given in the first argument.
`sysval`	Gives the exit value of the most recent system command.
`traceoff`	Turns the macro trace off.
`traceon`	Turns the macro trace on.
`translit`	Performs character transliteration.
`undefine`	Removes user-defined or built-in macro definitions.
`undivert`	Discards the diverted text.
`unix`	Null; indicates that the underlying system is derived from the UNIX operating system.

5 `lex`: A Lexical Analyzer

Overview of `lex` usage / 5-3
`lex` and `yacc` / 5-4
Program syntax / 5-6
Character set / 5-7
Definitions / 5-10
Rules / 5-12
Actions / 5-19
Compilation / 5-27
Examples / 5-27
Summary / 5-29

`lex` is a program generator that produces a program in a general-purpose language that recognizes regular expressions. It is designed for lexical processing of character input streams. It accepts high-level, problem-oriented specifications for character string matching.

Input to `lex` is a table of regular expressions and corresponding program fragments. The table is translated to a program that reads an input stream, copies the input stream to an output stream, and partitions the input into strings that match the given expressions. As each such string is recognized, the corresponding program fragment is executed.

The recognition of the regular expressions is performed by a deterministic finite automaton generated by `lex`. The program fragments are executed in the order in which the corresponding regular expressions occur in the input stream.

The code written by `lex` is not itself a complete language, but rather a generator representing a new language feature that can be added to different programming languages, called *host languages*. For example, one high-level language can be used for recognizing patterns, while a more general-purpose language is used for action statements.

The `lex` program generator can be used alone for simple transformations or for analysis and statistics gathering on a lexical level. The `lex` generator also can be used with a parser generator (for example, `yacc`) to perform the lexical analysis phase.

Just as general-purpose languages can produce code to run on different computer hardware, `lex` can write code in different host languages. The host language is used for the output code generated by `lex` and the program fragments that comprise the `lex` source program.

Compatible run-time libraries for the different host languages are provided, making `lex` adaptable to many environments and users. However, at present, the only supported host language is the C language.

Overview of `lex` usage

The program generated by `lex` is called `yylex`. The `yylex` program recognizes expressions in an input stream and performs the specified actions for each expression as it is detected. See Figure 5-1.

For example,

```
%%
[ \t]+$   ;
```

This sample `lex` source program is all that is required to generate a program to delete all blanks or tabs at the ends of the input lines. The `%%` delimiter is a `lex` convention to mark the beginning of the rules, the pattern-matching expressions. The rule itself,

```
[ \t]+$   ;
```

matches one or more instances of the characters blank and tab. The brackets enclose the character class consisting of blank and tab; the `+` indicates "one or more instance of the previous characters or character class" and the `$` indicates end-of-line. No action is specified, so the `yylex()` program (generated by `lex`) ignores these characters. Everything else is copied.

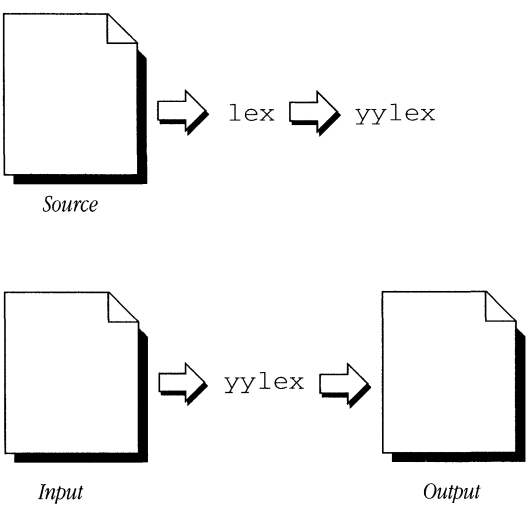

Figure 5-1 Overview of `lex`

Consider this next example:

```
%%
[ \t]+$   ;
[ \t]+    printf(" ");
```

The coded instructions in `yylex` scan for both rules at once. Once a string of blanks or tabs is recognized, `yylex` determines whether the string is followed by a newline character. If it is, then the first rule has been matched so that the corresponding action is performed; `yylex` does not copy the string to output. The second rule matches strings of one or more blanks and tabs not already satisfying the first rule, and causes `yylex` to replace a string of one or more blanks and tabs with a single space.

In `yylex`, the program generated by `lex`, the actions to be performed as each regular expression is found are gathered as cases of a switch. The automaton interpreter directs the control flow. It is possible to insert either declarations or additional statements in the routine containing the actions and to add subroutines outside this action routine, should you need to do so.

The `lex` program generator is not limited to one-character look-ahead. For example, if there are two rules, one looking for `ab` and another for `abcdefg`, and the input stream is `abcdefh`, `lex` recognizes `ab` and leaves the input pointer just before `cdefh`.

`lex` and `yacc`

It is particularly easy to use `lex` and `yacc` together. The `lex` program recognizes only regular expressions; `yacc` writes parsers that accept a large class of context-free grammars but requires a lower level analyzer to recognize input tokens. Thus, a combination of `lex` and `yacc` is often appropriate. When used as a preprocessor for a later parser generator, `lex` is used to partition the input stream; the parser generator assigns structure to the resulting pieces. The flow of control in such a case is shown in Figure 5-2. Additional programs, written by other generators or by hand, can be added easily to programs written by `lex`. The name "`yylex`" is what `yacc` expects its lexical analyzer to be named. If `lex` uses this name, it simplifies interfacing.

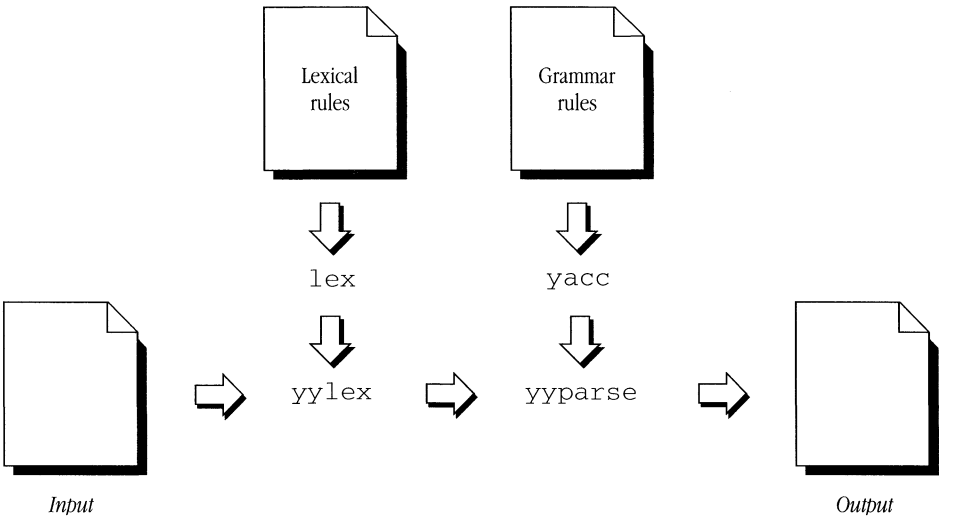

Figure 5-2 `lex` with `yacc`

To use `lex` with `yacc`, observe that `lex` writes a function named `yylex`, which is the name required by `yacc` for its analyzer. Normally, the default main program on the `lex` library calls the `yylex` routine, but if `yacc` is loaded and its main program is used, `yacc` calls `yylex`. In this case, each `lex` rule ends with

`return(`*token*`);`

where the appropriate token value is returned. An easy way to gain access to the names for tokens in `yacc` is to compile the `lex` output file as part of the `yacc` output file by placing the line

```
#include "lex.yy.c"
```

in the last section of the `yacc` input. If the grammar is to be named `good` and the lexical rules are to be named `better`, the command sequence could be

```
yacc good
lex better
cc y.tab.c -ly -ll
```

The `yacc` library (`-ly`) should be loaded before the `lex` library to obtain a main program that invokes the `yacc` parser. The generations of `lex` and `yacc` programs can be done in either order.

Program syntax

The general format of `lex` input is

```
{ definitions }
%%
{ rules }
%%
{ user subroutines }
```

where the *definitions* and the *user subroutines* are often omitted. The first `%%` is required to mark the beginning of the *rules*, but the second `%%` is optional. The absolute minimum `lex` program is

```
%%
```

This `lex` source generates a program that copies the input to output unchanged.

In the `lex` program format just shown, the *rules* consist of two parts:

- a left column with regular expressions
- a right column with actions and program fragments to be executed when the expressions in the left column are recognized

For example,

```
integer printf("found keyword INT");
```

The sample rule mentioned earlier gives the instructions to look for the string `integer` and, when found, produces the statement

```
found keyword INT
```

In this example, because the host procedural language is C, the C language library function `printf` is used to print the string.

The end of the expression is indicated by the first blank or tab character. If the action is a single C language expression, it can just be given in the right column, as illustrated in the example. If the action is compound or requires more than one line, it should be enclosed in braces. Consider the following example:

```
colour      printf("color");
mechanise   printf("mechanize");
petrol      printf("gas");
```

This `lex` source segment could be used to generate a program to change a number of words from British to American spelling. It should be noted, however, that these rules would have to be changed somewhat to be really useful. For example, if the word `petroleum` appeared in the input stream, the program generated by this segment would change it to `gaseum`.

Character set

Internally, a character is represented as a small integer. If the standard library is used, the value of a character is equal to the integer value of the bit pattern representing the character on the host computer. For example, the character `A` has the value \101 (octal) in ASCII.

Of course, you need not use the integer value of a character to access the value. The character `a` is represented in the same form as the character constant `'a'`. If this interpretation is changed by providing I/O routines that translate the characters, `lex` must be given a translation table that is in the *definitions* section of the source, and this translation table must be bracketed by lines containing only `%T`. The translation table, then, contains lines of the form

```
%T
{ integer } { character string }
%T
```

which indicate the value associated with each character.

Character classes

Classes of characters can be specified using the operator pair `[` and `]`. For example, the construction `[abc]` matches a single character, which can be `a`, `b`, or `c`.

Within brackets, most operator meanings are ignored. Only three characters are special:

```
-
\
^
```

The - character indicates a range. For example,

```
[a-z0-9<>_]
```

specifies the character class containing all the lowercase letters (a to z), digits (0 through 9), angle brackets (< and >), and the underline character (_).

Using - between any pair of characters that are not both uppercase letters, both lowercase letters, or both digits is sometimes acceptable to lex, but this is implementation-dependent. (It works on A/UX, but it might not be portable to other systems.) Therefore, if such a range is declared, lex issues a warning message. One reason for this is that [0-z] matches many more characters in ASCII than in EBCDIC.

If it is necessary to include the character - in a character class, it should either be first or last within the brackets. For example,

```
[-+0-9]
```

matches *all* digits (0 through 9) and the two symbols - and +.

The \ character acts as an escape character within class brackets. For example,

```
[a-z\*]
```

matches all lowercase letters (a to z) *and* the character *.

If the ^ operator appears as the first character after the left bracket, lex *ignores* the characters within the brackets, therefore matching all characters *except* those within the designated character class range. If an operation is to be performed on recognition of a string expressed using this construction, it is done on strings *other than* those within the brackets. For example,

```
[^abc]
```

matches all characters *except* a, b, or c, including all special and control characters. Also,

```
[^a-zA-Z]
```

matches any character that is *not* a letter (neither in the range a through z nor in the range A through Z).

Arbitrary characters

There are several other ways to specify characters to `lex`. The period operator (`.`) instructs `lex` to match any character except a newline. The meaning of the period does not change within brackets.

Also, all characters and ranges can be designated using the octal representations of those characters. This method, however, is difficult to read and most likely not portable. Nonetheless, the character class range

```
[\40-\176]
```

can be used to match all printable ASCII characters from octal 40 (blank) to octal 176 (tilde:~).

Operators

The operator characters are

```
" \ [ ] ^ - ? . * + | ( ) $ / { } % < >
```

If these are to be used as text characters, an appropriate "escape" should be used. For example, to get the character `\`, you must escape its significance as an operator. You can do so easily with another backslash: `\\`. For more information on escaping, refer to *A/UX Shells and Shell Programming.*

The quotation mark operator (`"`) indicates that whatever characters follow, up to a second `"` character, are to be taken as text characters without any "magic" meaning or operator significance. The quotation mark, then, is another way to escape the special meaning of a character. For example,

```
xyz"++"
```

matches the string `xyz++` wherever it appears. Of course, it is unnecessary, though harmless, to quote an ordinary text character. Consequently, the expression

```
"xyz++"
```

is equivalent to the one that quoted only the `++`. However, by quoting every character being used as a text character, you can avoid remembering the list of current operator characters, and avoid problems should further extensions to `lex` lengthen the list.

Another use of the quoting mechanism is for forcing a blank into an expression. Normally, as explained earlier, blanks or tabs end a rule. Any blank character not contained within brackets *must* be quoted.

There is also a third way to match the literal value of these operators, using the `\` escape character. You could specify the string discussed earlier as

```
xyz\+\+
```

Several C language escapes using `\` are recognized:

`\n`	newline
`\t`	tab
`\b`	backspace
`\\`	backslash

Since newline is illegal in an expression, `\n` must be used.

Definitions

Recall that the basic format of a `lex` source is

```
{ definitions }
%%
{ rules }
%%
{ user subroutines }
```

In addition to the *rules* (discussed later), `lex` includes options to define variables. Variables can occur either in the *definitions* section or in the *rules* section.

Remember, `lex` is generating the rules into a program, and any source not intercepted by `lex` is copied into the generated program. Also,

- Any line not part of a `lex` rule or action and that begins with a blank or tab is copied into the `lex` generated program.
- Any line not part of a `lex` rule or action that begins with a blank or tab and is found prior to the first `%%` delimiter is "external" to any function in the code.
- Any line not part of a `lex` rule or action that begins with a blank or tab and is found immediately after the first `%%` appears in an appropriate place for declarations in the function written by `lex` that contains the actions. This material must look like program fragments and should precede the first `lex` rule.

- Lines that begin with a blank or tab, and that contain a comment, are passed through to the generated program. This can be used to include comments in either the `lex` source or the generated code. The comments should follow the host language convention.
- Anything included *between* lines containing only `%{` and `%}` is copied to output. The delimiters are discarded. This format permits entering text-like preprocessor statements that must begin in column 1, or copying lines that do not look like programs.
- Anything after the third `%%` delimiter, regardless of formats, and so on, is copied to output *after* the `lex` output.

Definitions intended for `lex` are given before the first `%%` delimiter. Any line in this section not contained between `%{` and `%}` and beginning in column 1 is assumed to define `lex` substitution strings. The format of such lines is

name translation

This facility enables the string given as *translation* to be associated with the *name*. The *name* and *translation* must be separated by at least one blank or tab, and the *name* must begin with a letter. The *translation* can be called by the `{name}` syntax in a rule. Using `{D}` for the digits and `{E}` for an exponent field, you might have

```
D                       [0-9]
E                       [DEde][-+]?{D}+
%%
{D}+                    printf("integer");
{D}+"."{D}*({E})?             |
{D}*"."{D}+({E})?             |
{D}+{E}                 printf("real");
```

This example abbreviates rules to recognize numbers. The first two rules for real numbers both require a decimal point and contain an optional exponent field. The first requires at least one digit before the decimal point (`{D}+"."{D}*({E})?`), and the second requires at least one digit after the decimal point (`{D}*"."{D}+({E})?`). To correctly handle the Fortran expression `35.EQ.I`, which does not contain a real number, a context-sensitive rule such as

```
[0-9]+/"."EQ  printf("integer");
```

could be used, in addition to the normal rule for integers (see "Context Sensitivity").

The *definitions* section also can contain other commands, including the selection of a host language, a character set table, a list of start conditions, or adjustments to the default size of arrays within `lex` itself for larger source programs.

Repetitions and definitions

The operators `{` and `}` specify either

- repetitions (if they enclose numbers)
- definition expansion (if they enclose a name)

For example,

```
{digit}
```

looks for a predefined string named `digit` and inserts it at that point in the expression. The definitions are given in the first part of the `lex` input, before the rules. On the other hand, the expression

```
a{1,5}
```

looks for one to five occurrences of `a`.

An initial `%` is not an ordinary character, but has a special meaning to `lex` as the separator for source program segments.

Rules

Regular expressions

The regular expressions in `lex` function just as do those in the A/UX text editors `vi`, `ed`, and so on. A regular expression specifies a set of strings to be matched. It contains "text characters," which match characters in the input stream, and "operator characters," which, together with those "text characters," express a string that is to be recognized before the action in the right column takes place.

Letters of the alphabet and digits are always text characters. For example,

```
integer
```

matches the string `integer` wherever it appears, and the expression

```
a57D
```

looks for the string `a57D`.

Optional expressions

The question mark (`?`) operator indicates that what immediately precedes it is an optional element of an expression. Thus,

```
ab?c
```

matches either `ac` or `abc`.

Repeated expressions

Repetitions of classes are indicated by the operators `*` and `+`. The expression

```
a*
```

matches zero or more consecutive `a` characters. The expression

```
a+
```

matches one or more instances of `a` characters. The expression

```
[a-z]+
```

matches all strings of lowercase letters. The expression

```
[A-Za-z][A-Za-z0-9]*
```

matches all alphanumeric strings that have a leading alphabetic character. This is a typical expression for recognizing identifiers in computer languages.

Alternation and grouping

The operator `|` indicates *alternation.* For example,

```
(ab|cd)
```

matches either `ab` or `cd`. The parentheses are used here for grouping only. They are not required in such a simple and clear-cut example, but are often used for clarity or to force correct interpretation of more complex expressions. For example,

```
(ab|cd+)?(ef)*
```

matches such strings as

```
abefef
efefef
cdef
cddd
```

but not

```
abc
abcd
abcdef
```

Context sensitivity

The `lex` program recognizes a small amount of surrounding context. The two simplest operators for this are `^` and `$`.

As in the A/UX text editors, if the first character of an expression is `^`, the expression is matched only if found at the beginning of a line, either after a newline character or at the beginning of the input stream. Do not confuse this with the use of the `^` operator within brackets, which instructs `lex` to match any character except those in the designated character class range. If you want to use `lex` to find occurrences of a particular range of characters, but only if they occur as the first character on a line, you must use the `^` operator on the *outside* of the brackets. For example, the expression

```
^[0-9]
```

matches lines whose first character is a digit, `0` through `9`. The expression

```
^[^0-9]
```

matches lines whose first character is not a digit `0` through `9`.

The operator `$` is matched only at the end of a line, immediately followed by newline. This operator is a special case of the `/` operator character, which indicates "trailing context." The expression

```
ab/cd
```

matches the string `ab` only if followed by `cd`. Therefore, the expression

```
ab$
```

can also be expressed

```
ab/\n
```

That is, the use of the `$` operator can be interpreted as an instruction to match the characters only when followed by a newline.

Left context is handled in `lex` by "start conditions." If a rule is only to be executed when the `lex` automaton interpreter is in start condition `x`, the rule should be enclosed within the angle-bracket operator characters:

```
<x>
```

If "being at the beginning of a line" is considered to be start condition `ONE`, then the `^` operator is equivalent to

```
<ONE>
```

See the sections entitled "Left Context Sensitivity," "Examples," and "Summary" for further explanation and illustration of start conditions.

Left context sensitivity

Sometimes it is desirable to have several sets of lexical rules applied at different times in the input. For example, a compiler preprocessor might distinguish preprocessor statements and analyze them differently from ordinary statements. This requires "sensitivity" to prior context. There are several ways of handling such occurrences. For example, the `^` operator is a "prior context operator" because it must recognize the immediately preceding left context to discern whether a character appears at the beginning of a line, just as the `$` operator must recognize the immediately following right context to discern whether a character appears at the end of a line.

Adjacent left context can be extended to produce a facility similar to that for adjacent right context. This is likely to be less useful, however, since often the relevant left context, such as the beginning of a line, appeared some time earlier.

There are three basic ways of dealing with different environments so as to achieve a lexical analysis with a greater degree of context sensitivity.

- The use of flags. This is most useful when only a few rules change from one environment to another.
- The use of start conditions on rules.
- The possibility of making multiple lexical analyzers all run together. If the sets of rules for the different environments are very dissimilar, clarity might best be achieved by writing several distinct lexical analyzers and switching from one to another, as necessary.

In each case, there are rules that recognize the need to change the environment in which the following input text is analyzed and a parameter is set to reflect the change. The remainder of this section describes in greater detail the first two ways of dealing with different environments.

Flags

The simplest way of changing the environment in which input is analyzed is by use of a *flag* explicitly tested by the user's action code. If done in this way, `lex` is not involved at all.

To illustrate, consider the following program requirements:

- Copy the input to the output.
- Change the word `magic` to `first` on every line that begins with the letter `a`.
- Change `magic` to `second` on every line that begins with the letter `b`.
- Change `magic` to `third` on every line that begins with the letter `c`.

All other words and all other lines are left unchanged. These rules are so simple that the easiest way to do this job is with a flag. For example,

```
        int flag.
%%
^a      {flag = 'a'; ECHO;}
^b      {flag = 'b'; ECHO;}
^c      {flag = 'c'; ECHO;}
\n      {flag = 0 ; ECHO;}
```

```
magic {
        switch (flag)
        {
          case 'a': printf("first"); break;
          case 'b': printf("second"); break;
          case 'c': printf("third"); break;
          default: ECHO; break;
        }
      }
```

Start conditions

It might be more convenient to have `lex` "remember" the flags as start conditions on the rules. Any rule can be associated with a start condition. That rule, then, is recognized only when `lex` is in that start condition. The current start condition can be changed at any time. To handle the same problem using start conditions, begin by introducing each start condition to `lex` in the *definitions* section with a line reading

`%Start` *name1 name2* ...

where the conditions (*name1*, *name2*, and so on) can be named in any order. The word `Start` can be abbreviated to `s` or `S`. Then, to reference the conditions, use angle brackets:

`<name1>` *expression*

The rule illustrated earlier is recognized *only* when `lex` is in the start condition *name1*. To enter that start condition, execute the following action statement:

`BEGIN` *name1*`;`

The action statement

`BEGIN 0;`

resets the initial condition of the `lex` automaton interpreter.

A rule can be active in several start conditions. For example,

`<`*name1*`,` *name2*`,` *name3*`>` *expression*

is a legal expression. Any rule *not* beginning with the `<` prefix operator is always active.

The following example illustrates the use of start conditions:

```
%START AA BB CC
%%
^a     {ECHO; BEGIN AA;}
^b     {ECHO; BEGIN BB;}
^c     {ECHO; BEGIN CC;}
\n     {ECHO; BEGIN 0;}
<AA>magic   printf("first");
<BB>magic   printf("second");
<CC>magic   printf("third");
```

Obviously, this example is a rewrite of the previous example; the problem-solving logic is exactly the same. However, in this case `lex` was instructed to do the work instead of the host language code.

Ambiguous rules

The `lex` program can handle ambiguous specifications. When more than one expression can match the current input, the longest match is preferred, among rules that matched the same number of characters, the rule given first is preferred. For example, using the rules

`integer` *keyword-action* ;
`[a-z]+` *identifier-action* ;

(if the input were `integers`), `lex` interprets the input as an identifier because `[a-z]+` matches all eight characters (including the final `s`), while `integer` matches only seven characters.

If the input were `integer`, both rules would match the seven characters. In that case, `lex` selects the keyword rule because it was given first. If the input were anything shorter (for example, `int`), the input would not match the expression `integer`. It would, however, match the `[a-z]+` expression, so the identifier interpretation would be used.

The principle of preferring the longest match makes rules containing expressions like `.*` dangerous. For example,

```
'.*'
```

appears to instruct `lex` to find a match for a string in single quotes. However, it is an instruction for the program to read far ahead looking for a distant single quote. For example, if the above expression were given the following input:

```
'first' quoted string here, 'second' here
```

the expression would match almost the entire input line:

```
'first' quoted string here, 'second'
```

which is most likely *not* the desired result. A better rule for matching strings within single quotes might be

```
'[^'\n]*'
```

which, given the same input, matches `'first'`. The consequences of errors like this are greatly lessened by the fact that the period (`.`) operator does not match newline. Expressions like `.*` stop on the current line.

◆ **Note** Do not try to defeat the protection of `.` not matching the newline character with expressions such as `[.\n]+` or an equivalent, because the program generated by `lex` then tries to read the entire input file, causing internal buffer overflows. ◆

Actions

When an expression written as the previous one is matched, `yylex` executes the corresponding action. The default action for `yylex` is to copy input to output, and is performed on all strings not otherwise matched. Therefore, a rule that merely copies can be omitted. If you want to absorb the entire input without producing any output, you must provide rules to match everything. (When `yylex` is being used with `yacc`, this is the normal situation.) In other words, by default, a character combination in input that was omitted from the rules is printed on the output.

The null statement

One of the simplest things that can be done is to ignore the input. To accomplish this, use a semicolon (`;`) as the action (a semicolon is the C language "null statement"). The rule

```
[ \t\n]   ;
```

causes the spacing characters (that is, blank, tab, and newline) to be ignored because it gives the null statement as its associated action.

The repetition character

The vertical bar character (`|`) represents the instruction to use the action designated for the next rule for the current rule as well. For example,

```
" "     |
"\t"    |
"\n"    ;
```

This example instructs `yylex` to ignore the spacing characters, as did the previous example. The first line gives the rule "match blank characters" and instructs the program to perform the action indicated for the next rule. Then, the second line gives the rule "match `\t` characters" and instructs the program to perform the action indicated for the next rule. Finally, the third line gives the rule "match `\n` characters," and gives the action `;`, the null statement. Therefore, the action for all three rules is the null statement.

`printf` and `ECHO`

In more complex actions, you might often want to know the actual text that matched a regular expression. The `yylex` program leaves this text in an external character array, named `yytext`. Consider the following example:

```
[a-z]+  printf("%s", yytext);
```

This example illustrates a way of accessing the characters matching a regular expression. Using this example, the rule given is to find the strings matching the regular expression `[a-z]+` and the action is to print those strings in the character array `yytext` using the C language function `printf`.

The `printf` function accepts a format argument and data to be printed. Still using this example, the format is `%s` (print string). The `%` character indicates data conversion, and `s` indicates data type string, in this case the character array `yytext`. This places the matched string on the output.

The action of printing the strings matching the regular expressions is so common that it can be written simply as `ECHO`. For example,

```
[a-z]+   ECHO;
```

This example accomplishes the same action as the previous one using the `printf` statement.

Even though the default action is to copy input to output, the `ECHO` facility is included explicitly to provide a more discriminating copy function. For example, a rule that matches `read` normally matches all instances of `read`, even those contained in other words (`bread`, `treadmill`, and so on). To avoid this, a rule of the form `[a-z]+` is needed. This is explained further in the following section.

yyleng

Sometimes it is necessary to know what is at the end of a matched pattern. To facilitate this, `lex` provides a count of the number of characters matched, `yyleng`. To count both the number of words in the input and the number of characters in those words, you might write

```
[a-zA-Z]+   {words++; chars += yyleng;}
```

This instruction takes the strings that match the regular expression `[a-zA-Z]+` and accumulates the number of characters in these strings in `chars`. Then, the action instruction

```
yytext[yyleng-1]
```

can be used to access the last character in the string matched.

yymore and yyless

Occasionally, a lex action might decide that a rule did not recognize the correct span of characters. Two routines are provided to aid with this situation:

yymore() This routine instructs yylex to tack the next input expression recognized on to the end of this input. Normally, the next input string overwrites the current entry in yytext.

yyless(*n*) This routine instructs yylex to retain in yytext only *n* (a number) of those characters resulting from the current expression. Further characters previously matched are returned to the input. This provides the same sort of look-ahead offered by the / operator, though in a very different form.

Consider a language that defines a string as a set of characters between quotation marks ("), and requires that the " character be preceded by a \ to be included in a string. The regular expression which matches that is somewhat confusing, so it might be preferable to write the following segment:

```
\"[^"]* {
        if (yytext[yyleng-1] == '\\')
                yymore();
        else
                ...normal user processing
}
```

The previous lex segment, when it finds the string

```
"abc\"def"
```

first matches the five characters "abc\ and then calls the yymore routine, which causes the next part of the string, "def, to be tacked on the end of the input. Note that the final quote terminating the string should be picked up in the code labeled *normal user processing.*

The function yyless might be used to reprocess text in various circumstances. Consider, for example, the problem of disambiguating a C language statement such as

```
s=-a
```

One way to parse this statement treats the `-` as part of the operator:

```
=-[a-zA-Z] {
            printf("Operator (=-) ambiguous\n");
            yyless(yyleng-1);
            action for =-
          }
```

This `lex` segment prints a message, treats the operator as `=-`, and returns the letter found after the operator to the input stream. However, you might want to treat this syntax as `= -a`. In that case,

```
=-[a-zA-Z] {
            printf("Operator (=-) ambiguous\n");
            yyless(yyleng-2);
            action for =
          }
```

prints a message, treats the operator as `=`, and returns `-a` to the input stream.

It is possible to avoid the misinterpretation of operators by rewriting the regular expression. To indicate that the operator is `=-`, using the same example, use the following rule:

```
=-/[A-Za-z]
```

To indicate that the operator is `=`, use the following rule:

```
=/-[A-Za-z]
```

No backup is required in the rule action. It is not necessary to recognize the whole identifier to observe the ambiguity. However, the possibility of `=-3` makes

```
=-/[^ \t\n]
```

a still better rule.

`lex` input and output routines

The programs generated by `lex` handle character I/O only through the routines `input`, `output`, and `unput`. The character representation provided in these routines is accepted by `lex` and used to return values in `yytext`. These are provided as `lex` macro definitions, as shown in the following list.

`input()`	returns the next input character
`output(c)`	writes the character *c* on the output
`unput(c)`	pushes the character *c* back onto the input stream to be read later by `input`

(As shown previously, you can use `printf` to generate error messages.) These routines are provided by default, but you can override them by providing your own versions. To redefine or override a `lex` routine, include your own version in the *user subroutines* section. These routines must be standard C and be named according to the `lex` routine you want to replace. However, because these routines define the relationship between external files and internal characters, they must all be retained and/or modified consistently.

These routines can be redefined to cause input or output to be transmitted to or from other programs or internal memory. The character set used must be consistent in all routines and a value of 0 returned by `input` must mean end-of-file.

The relationship between `unput` and `input` must be retained or the `lex` look-ahead does not work. The `lex` program does not look ahead at all if it does not have to; rules ending in `+`, `*`, `?`, or `$`, or those containing a `/`, however, force look-ahead. Look-ahead is necessary to match an expression that is a prefix of another expression. The standard `lex` library imposes a 100-character limit on backup.

yywrap

Another `lex` library routine that you might sometimes want to redefine is `yywrap`. To redefine or override a `lex` routine, include your own version in the *user subroutines* section. These routines must be standard C and be named according to the `lex` routine you want to replace. This routine is called whenever `lex` reaches an end-of-file. If `yywrap` returns a 1, which it does by default, `lex` continues with the normal wrap-up on end of input.

It is sometimes convenient to arrange for `input` to continue from a new source. In this case, `yywrap` can be redefined to arrange for new input and return 0. This then instructs `lex` to continue processing.

This routine provides a convenient way to print tables, summaries, and so on, at the end of a program. It is not possible to write a normal rule that recognizes end-of-file. The only access to this condition is through `yywrap`. In fact, unless a private version of `input` is supplied, a file containing nulls cannot be handled because a value of 0 returned by `input` is taken to be end-of-file by `yywrap`.

REJECT

Note that `lex` is normally partitioning the input stream, not searching for all possible matches of each expression. This means that each character is accounted for once and only once. Consider the following example:

```
she     s++;
he      h++;
\n      |
.       ;
```

The first rule matches all occurrences of the string `she` and the action increments `s` for each one found. The second matches all occurrences of the string `he` and its action increments `h` for each one found. The last two rules match newline and everything else and take the action of ignoring them. Normally, `lex` would not recognize the instances of `he` included in `she`, because once it passed a `she`, those characters are gone. To override this default, the action `REJECT` can be used to instruct `lex` to go to the next alternative. `REJECT` causes the rule *after* the current rule to be executed. The position of the input pointer is adjusted accordingly.

Suppose you want to count the instances of `he` included in `she`. To do that, use the following rules:

```
she     {s++; REJECT;}
he      {h++; REJECT;}
\n      |
.       ;
```

In this example, after counting each expression, the expression is "rejected" (whenever appropriate), and the other expression is evaluated. In this example, because `he` does not include `she`, the `REJECT` action on `he` can be eliminated. In other cases, it is not possible to state which input characters are in both classes.

Consider the following two rules:

```
a[bc]+    { ... ; REJECT;}
a[cd]+    { ... ; REJECT;}
```

- If the input to the rules above were `ab`, only the first rule would match.
- If the input to these same rules were `ad`, only the second would match.
- If the input were `accb`, the first rule would match four characters and the second rule would match three characters.
- If the input were `accd`, however, the second rule would match four characters and the first rule would match three characters.

In general, `REJECT` is useful whenever the purpose of `lex` is to detect all examples of some items in the input for which the instances of these items might overlap or include one another, instead of the usual purpose of `lex` of partitioning the input stream.

Suppose you want a diagram of some input. Normally, the digrams overlap, that is, the word `the` is considered to contain both `th` and `he`. Assuming a two-dimensional array named `digram[]` to be incremented, an appropriate `lex` procedure is

```
%%
[a-z][a-z] {digram[yytext[0]][yytext[1]]++; REJECT;}
.          |
\n         ;
```

In this example, `REJECT` is used to pick up a letter pair beginning at every character, rather than at every other character.

The action `REJECT` does not rescan the input. Instead, it "remembers" the results of the previous scan. Therefore, if `yylex` is instructed to find a rule with trailing context and execute `REJECT`, `unput` cannot have been called to change the characters forthcoming from the input stream. This is the only restriction on the user's ability to manipulate the not-yet-processed input.

Compilation

The following steps are involved in compiling a `lex` source file:

1. The `lex` source must be transformed into a program in the host general-purpose language. The generated program is put into a file named `lex.yy.c`.
2. That program must then be compiled and loaded, usually with a library of `lex` subroutines. The I/O library is defined in terms of the C language standard library. On the A/UX operating system, the library is accessed by the loader flag `-ll`. In this case, an appropriate set of commands is

   ```
   lex inputfile
   cc lex.yy.c -ll
   ```

   The resulting program is placed in the file `a.out` for later execution.

Although the default `lex` I/O routines use the C language standard library, `lex` routines such as `input`, `output`, and `unput` do not. Therefore, if your own versions of these routines are given, the library is avoided.

Examples

For the sake of example, consider copying an input file while adding three to every positive number divisible by 7. A suitable `lex` source program follows:

```
%%
   int k;
[0-9]+  {
          k = atoi(yytext);
          if (k%7 == 0)
            printf("%d", k+3);
          else
            printf("%d", k);
        }
```

The rule `[0-9]+` recognizes strings of digits, 0 through 9; `atoi` converts the digits to binary and stores the result in `k`. The operator `%` (remainder) is used to check whether `k` is divisible by seven; if it is, `k` is incremented by 3 as it is written out. It might be objected that this program alters such input items as `49.63` or `X7`. Furthermore, it increments the absolute value of all negative numbers divisible by 7. To avoid this, add a few more rules after the active one. For example,

```
%%
   int k;
-?[0-9]+  {
            k = atoi(yytext);
            printf("%d", k%7 == 0 ? k+3 : k);
          }
-?[0-9.]+             ECHO;
[A-Za-z][A-Za-z0-9]+ ECHO;
```

Numeric strings containing a period (`.`), or preceded by a letter, are picked up by one of the last two rules and not changed. The `if-else` is replaced by a C language conditional expression to save space. The expression *a* `?` *b* `:` *c* is evaluated as "if *a* then *b* else *c*."

The following is an example using `lex` for gathering statistics. This program reports how many words of various lengths there are. (A word is defined here as a string of letters.)

```
         int lengs[100];
%%
[a-z]+   lengs[yyleng]++;
.        |
\n       ;
%%
yywrap( )
{
int i;
printf("Length  No. words\n");
for(i=0; i<100; i++)
   if (lengs[i] > 0)
      printf("%5d%10d\n", i, lengs[i]);
```

```
return(1);
}
```

In the preceding example, the data is accumulated but no output is generated until, at the end of the input, the table is printed. The final statement, `return(1);`, indicates that `lex` is to perform wrap-up. If `yywrap` returns 0 (false), it implies that further input is available and the program is to continue reading and processing. Remember, providing a `yywrap` that never returns true causes an infinite loop.

Summary

The general form of a `lex` source file is

```
{ definitions }
%%
{ rules }
%%
{ user subroutines }
```

The *definitions* section contains a combination of the following items:

- Definitions in the form

  name translation

- Included code in the form

  code

  where a space (or tab) must precede *code.*

- Included code in the form

  `%{`

  code

  `%}`

- Start conditions given in the form

  `%S` *name1 name2* ...

- Character set tables in the form

  `%T`

  number character-string

  ...

  `%T`

- Changes to internal array sizes in the form

  `%x` *nnn*

  where *nnn* is a decimal integer representing an array size and `x` selects the parameter as follows:

Letter	*Parameter*
`p`	positions
`n`	states
`e`	tree nodes
`a`	transitions
`k`	packed character classes
`o`	output array size

Lines in the *rules* section have the form

expression action

where the *action* can be continued on succeeding lines by using braces to delimit it.

Regular expressions in `lex` use the operators shown in Table 5-1.

Table 5-1 Regular expression operators

Expression	Meaning
`x`	The character `x`
`"x"`	An `x`, even if it is an operator
`\x`	An `x`, even if it is an operator
`[xy]`	The character `x` or `y`
`[x-z]`	The characters `x`, `y`, or `z`
`[^x]`	Any character but `x`
`.`	Any character but newline
`^x`	An `x` at the beginning of a line
`<y>x`	An `x` when `lex` is in start condition `y`
`x$`	An `x` at the end of a line
`x?`	An optional `x`
`x*`	0 or more instances of `x`
`x+`	1 or more instances of `x`
`x|y`	An `x` or a `y`
`(x)`	An `x`
`x/y`	An `x`, but only if followed by `y`
`{xx}`	Expands to `xx` definition in `lex` definition section
`x{`*m*`,`*n*`}`	*m* through *n* occurrences of `x`

Part 2 File Manipulation Tools

The A/UX tools detailed in this section help you perform file-related tasks such as finding a file size or location, determining the differences between two files, and obtaining the version number of a program. Additionally, A/UX provides tools to control the file versions to ensure that they are the most recent and provides a way of updating and maintaining groups of files.

The chapter "File Attribute Tools" describes the tools to

- compare source files: `diff` and `comm`
- find files: `find`
- determine file characteristics: `size`
- find the version number of a file: `version`
- maintain portable archives: `ar`

The following chapters in this section describe the file maintenance tools to

- maintain and keep track of related program files: `make`
- manage versions of source code: `SCCS`
- process and parse files: `awk`

6 File Attribute Tools

Comparing source files / 6-2
Finding files: `find` / 6-2
Printing the section sizes of COFF files: `size` / 6-2
Finding the version number of a file: `version` / 6-3
Maintaining portable archives and libraries: `ar` / 6-3

The A/UX tools detailed in this section help you perform file-related tasks such as finding a file size or location, determining the differences between two files, and obtaining the version number of a program.

Comparing source files

A/UX includes a number of programs that compare files, including

`bdiff`	Used similarly to `diff`; its purpose is to allow processing of files that are too large for `diff`.
`diff`	A differential file comparator. It tells what lines differ in two files.
`diff3`	A three-way differential file comparator, which works only on files less than or equal to 64K bytes. It compares three versions of a file and publishes disagreeing ranges of text, flagged with special codes.
`diffmk`	Marks the differences between files. It compares two versions of a file and creates a third file that includes "change mark" commands for the `nroff` and `troff` formatters.
`diffdir`	Compares the differences in two directories of files.
`comm`	Selects or rejects lines common to two sorted files.

Finding files: `find`

`find` is a powerful utility that performs a depth-first recursive search for files of a given characteristic such as name, group, owner name, time of last modification or access, and so on. See `find`(1) in *A/UX Command Reference* for more information.

Printing the section sizes of COFF files: `size`

The `size` command produces size information for common object format files (COFF). See `size`(1) in *A/UX Command Reference* for more information.

Finding the version number of a file: `version`

`version` is useful for determining which version of a program you are running. `version` takes a list of files and reports the version number for each. If the file is not a binary, it reports that. If a version number is not associated with the file, the program reports that fact. `version` also reports the object file format of each file—that is, either `Coff object file format`, or `Old a.out object file format`.

The user can associate a version number with a file by defining a string constant at the top of the source code. The string must have the form

```
{Apple version RELEASE.LEVEL YY/MM/DD HH:MM:SS}"
```

In this string, the words *Apple version* must appear followed by the values for the release number, level number, year, month, day, hour, minute, and second. For example

```
char *_Version_ =  \
  "(c) Copyright 1986  {Apple version 2.1 86/09/12 18:05:24}"
```

See `version`(1) in *A/UX Command Reference* for more details.

Maintaining portable archives and libraries: `ar`

You can use the archive command `ar` to combine several files into one archive. An **archive** consists of a collection of files, plus a table of contents. Archives are used mainly as libraries to be searched by the link editor `ld`. A **library** (or library archive) is an archive that contains object files (plus a table of contents). Putting together your own library allows you to use locally produced functions (instead of limiting you to the functions supplied in standard libraries).

`ar` also provides the facility to append files to and delete files from the archive. Because the order of files is so important to the efficient operation of `ld`, you can also move files around within the archive, as well as extract them, print them, and produce a table of contents. See `ar`(1) in *A/UX Command Reference* for more information.

7 `make`: A File Production Tool

Using `make` / 7-3
The description file / 7-8
Suffixes and rules / 7-20
Operation / 7-28
Advanced topics / 7-44

The `make` program automates the production of related sets of files. It simplifies the task of administering libraries, functions, related source and object files, and many other administration tasks that must reflect a change when you update one file in the set. Although `make` is normally used to maintain program code, it can also be used for other batch data-processing activities. (For example, `make` is often used to produce technical manuals with `troff`.)

The `make` program keeps track of file dependencies; when you change one part of a program, `make` recompiles related files with a minimum amount of effort. The required information is maintained by the `make` program itself (which has built-in "rules" for recompilation), by using certain system information, such as the time stamp of the files, and by the description of operations kept in a file called the *description file* or *makefile.* Once you set up a makefile for a large project, `make` keeps track of your dependencies for you and frees you to concentrate on programming or other tasks.

Using `make`

The simplest use of `make` is

```
make file1
```

where a file named `file1.c` resides in the current directory. The file `file1.c` can use information from other files by using `#include` statements. This command causes `make` to find `file1.c` in the local directory and issue the proper command to compile it into `file1`.

◆ **Note** If `file1.c` has the same filename prefix (the same filename without the `.c` suffix) as another file, `make` might compile that file instead. If, for example, there is a more recent `file1.l` file, it is compiled instead, and `file1.c` is overwritten in the process. If these files are not different incarnations of the same program, losing the `.c` file could be quite troublesome. ◆

As long as only one file is involved and only a standard compilation is required, you do not need to create a makefile to `make` your files.

If, however, your program is spread over multiple files, you do need to create a **makefile,** which is a control file containing the filenames, a description of their interrelations, and actions to be performed on them. When it does not have enough to go on, `make` looks in the current directory for a file named `makefile` (or `Makefile`) that contains the necessary administrative information. In general, you must put an entry in the makefile for any file that has a nonstandard compilation procedure.

Writing a makefile

To write a makefile, you must determine the following elements:

- the target filename
- filenames of related compilation units (files)
- file dependencies
- related libraries
- the command that produces the target (including options for the programs to be run)

Targets are filenames, or placeholders for filenames, that are meant to be compiled.

The `make` program defines a **dependency** as follows: *file1* depends on *file2* only if *file1* needs to be recompiled whenever *file2* is changed. For example, if file `x.c` contains the line

```
#include "defs.h"
```

the object file `x.o` depends on `defs.h`. If `defs.h` is changed, the `x.o` file must be remade by compiling `x.c`. Note that the `x.c` (source) file does not depend on `defs.h`, because it does not need to be re-created when `defs.h` changes.

For example, a file named `zeke` depends on `zeke.o` and uses library functions from `libm.a`. To relink `zeke`, you enter

```
cc -lm zeke.o -o zeke
```

`cc` passes the two options, `-lm` and `-o`, to `ld`.

- `-lm` causes library `libm.a` to be searched.
- `-o` renames the compiled binary file `zeke` (instead of the default `a.out`).

The example requires the following makefile:

```
zeke: zeke.o
(TAB)     cc -lm zeke.o -o zeke
```

The first line states the dependency (that `zeke` depends on `zeke.o`). The second line is the command line describing the action that must take place whenever `zeke.o` changes. The command line must begin with a tab (represented by `(TAB)` in the examples).

In a more complicated example, a file named `xavier` depends on three files named `yancy.o`, `quincy.o`, and `wally.o`, all of which depend on `defs.h` and use the library `libm.a`. The command to link `xavier` is

```
cc -o xavier yancy.o quincy.o wally.o -lm
```

The makefile for `xavier` follows:

```
xavier: yancy.o quincy.o wally.o
(TAB)     cc yancy.o quincy.o wally.o -o xavier -lm
yancy.o quincy.o wally.o: defs.h
```

When makefiles become more complicated, you can use macros and other features described in the sections that follow.

When you have included the interfile dependencies and command sequences in a makefile, the command

```
make
```

updates the appropriate files, regardless of how many files you edited since the last time you executed `make`. The `make` utility uses the date and time that a file was last modified to find files that are out of date with respect to their targets.

`make` command syntax

`make` uses the following command syntax:

`make` [*option* ...] [*macro*=*def*...] [`-f` *filename*] [*target* ...]

`make` interprets these arguments in the following order:

1. First, `make` analyzes the macro definition arguments (arguments with embedded equal signs) and the assignments made. Command-line macros override corresponding definitions found in the description files. See "Macro Definitions" for more information.
2. Next, `make` examines the options. See "Options" for details.
3. Finally, `make` assumes the remaining arguments to be the names of targets to be made, and these are made in the order in which they appear on the command line. If there are no remaining arguments, the first target in the description file is made.

◆ **Note** `make` finds the first target by scanning the description file for a target that does not represent an internal file transformation rule (see "Transformation Rules"). These "built-in" rules are of the form

```
.n[.m]:
```

Where `n` and `m` are suffixes, any rule begins with a period and contains no slashes (as a full pathname might). Thus, the first target is the first name in the description file that does not begin with a period or begins with a period but contains a slash. ◆

Options

`make` accepts the following options:

`-a`	Update all targets. This option is useful for completely rebuilding all files.
`-b`	Use compatibility mode for old makefiles. This mode is on by default.
`-B`	Turn off compatability mode.
`-d` *digits*	Debug mode. If specified without *digits*, full debug mode is invoked. If specified with *digits*, a particular level of debugging is invoked. Debug levels 0, 1, and 2 tell you increasing levels of information about the `make` operation. Level 4 shows you how the macros are expanded. Level 9 displays actual flag names. The `-d` option also can be invoked by sending the signal `USR1` to `make`.
`-e`	Cause environment variables to override macro definitions.
`-f` *filename*	Use a different description file. *filename* is the name of a description file. A *filename* of hyphen (`-`) denotes the standard input. If there are no `-f` arguments, `make` reads the file named `makefile` or `Makefile` in the current directory in the order stated. Failing these, `s.makefile` or `s.Makefile` is sought in the SCCS directory, if such a directory exists. If a description file is present, its contents override the default rules.
`-g`	Turn on capabilites to automatically check out SCCS files. See "SCCS File Handling."
`-G`	Enable the Dynamic Include File Dependency Generation (DIFDG). The DIFDG also can be enabled by defining the variable `MAKEDIFDGSUFFIXES` as a list of legal suffixes for the source files to be searched.
`-i`	Ignore error codes that might be returned by a shell command. This mode can be entered if the target name `.IGNORE:` appears in the description file. See "Built-in Targets."
`-k`	If a shell command returns a nonzero status, abandon work on the current target but continue to process other targets that do not depend upon the abandoned target. (Targets are described under "Makefile Entries," and branches are discussed in "The `make` Predecessor Tree.")

`-K` Turn off the `-k` option. The `-K` option is on by default. The `-K` option is most often used in a descripton file that invokes `make`, that is a member of a multilevel `make` hierarchy, and that is invoked by a top-level `make` with the `-k` option.

`-M` Store the dependency map for all object files in a file. The default name for this file is `._Make_State`. (You can change the default name by changing the value in the `MAKEDEPFILE` variable to the desired name.) This option also can be enabled by defining the variable `MAKEDEPFILE` as the name of the file in which you want to store the map.

`-n` No-execute mode. Print the commands in the description file as they would be executed, but do not actually execute them. Even lines beginning with an `@` (at) sign are printed. However, if a command line has the string `$(MAKE)` in it, the line is always executed.

`-p` Print out the built-in rules of make including a complete set of macro definitions.

`-P` Search for `PRE` and `POST` files in the directory `/usr/lib`. For example, for a description file named `x.mk`, `make` searches for and reads `/usr/lib/x.mkPre` and `/usr/lib/x.mkPost`.

`-q` Question mode. The `make` command returns a zero or nonzero status code, depending on whether the target file is up to date.

`-r` Do not use the built-in rules of `make`. To do any useful work, this option must be accompanied by an appropriate description file.

`-s` Silent mode. Do not print command lines before executing them. This mode is entered if the built-in target name `.SILENT:` appears anywhere in the description file.

`-t` Update the target files using the `touch` command without executing any commands in the target files.

`-u` Look for `makecomm` and `Makecomm` files in the user's home directory, as specified by the `$HOME` environment variable, and in the current directory. The search order is `$HOME/makecomm`, `$HOME/Makecomm`, `./makecomm` and `./Makecomm`. At most, `make` reads one file from each directory. These files are read before any description files and can be used to define macros and rules.

`-V` Display current version of `make`.

Using `make` on individual files

Individual files mentioned in the makefile also can be used as arguments on the command line, if you want to compile only a single file. For example, with the makefile from the previous example

```
xavier: yancy.o quincy.o wally.o
(TAB)     cc  yancy.o quincy.o wally.o -o xavier -lm
yancy.o quincy.o wally.o: defs.h
```

and the command line

```
make yancy.o
```

`make` remakes only `yancy.o`, including `defs.h` in the process. To make both `yancy.o` and `wally.o`, you type

```
make yancy.o wally.o
```

and both files are remade properly.

The description file

The description file (often called the makefile) defines the target file and its dependencies. A description file can contain the following elements:

- makefile entries, consisting of dependency statements and commands or command sequences
- comments
- `include` lines
- macro definitions

◆ **Note** If you do not supply a description file, `make` uses its default rules to produce the file named on the command line. See the section "The Default Rules," later in this chapter. If you name your description file something other than `makefile` or `Makefile`, you must use the `-f` option on the `make` command line. See the section "Options," later in this chapter, for details. ◆

Makefile entries

A **makefile entry** defines the relationship between a target and its dependents and usually stipulates the command as well. A description file often contains multiple entries. The general form of a makefile entry is

target1 [*target2* ...] `:` [`:`] [*dependent1*...] [`;` *commands*] [`#` *comment*]
[(TAB) *commands*][#...]
[(TAB) *commands*][#...]
...

where (TAB) represents a tab character. Shell metacharacters such as `*` and `?` are expanded only in the command sequence. For example,

```
zeke: zeke.o
(TAB)     cc zeke.o -o zeke -lm
```

◆ **Note** Even though the tables generated by the makefiles are dynamically allocated, there are certain limits for the length of a line and the number of targets per line. If you run into problems with these limits, you can use the `adb` debugger to increase the `max_GPBuffer_size` for line length and `max_lefts_entries` for the number of targets per line. For more information on the `adb` debugger, see *A/UX Programming Languages and Tools*, Volume 1. ◆

Targets versus rules

Within a description file, user-defined rules can replace the built-in rules of `make`. User-defined rules can appear in the makefile entry anywhere a target name can be given.

You also can create a `/usr/lib/MakeRules` file that overrides the built-in rules. This allows site-specific rules. In either case, `make` still observes the `-r` option.

Some aspects of rule syntax are similar to target syntax. A target can be differentiated from a rule by the following criteria:

- A target name can begin with or without a period, and it contains slashes.
- A rule begins with a period and does not contain slashes. (See "Transformation Rules" for more information.)

Built-in targets

Not all targets correspond to files. `make` has defined certain **built-in targets** (targets to which no files correspond) to modify the behavior of `make`. These targets are passed to `make` in the description file. Because `make` reads the entire description file before beginning to process dependency statements, the built-ins, which must appear at the beginning of a line, are processed first, whether they appear at the beginning, middle, or end of the description file. Examples of built-in targets are as follows:

`.DEFAULT:`	If a file must be made but there are no explicit shell commands or relevant built-in rules, `make` uses the shell commands listed under `.DEFAULT:`.
`.IGNORE:`	If present, `.IGNORE` has the same effect as the `-i` option, which is to ignore nonzero return codes from commands.
`.PRECIOUS:`	The default behavior of `make` is to remove a target and its dependents when a quit or interrupt signal is received while processing the commands that update the target. Because the actions of `make` depend in large part on the mere existence of a file, removal of potentially incomplete files helps ensure that the proper files are regenerated each time. Removal can be avoided by making specific files dependent on `.PRECIOUS:`.
`.MAKESTOP [` *exit-code* `]:`	If present, `.MAKESTOP:` causes `make` to exit. `.MAKESTOP:` is useful in a multilevel directory and description file hierarchy to quickly bypass a `make` in one particular directory or in several directories. *exit-code* is optional and defaults to zero if not specified. If *exit-code* is not specified or if the specified exit code is zero, `make` exits silently. If a nonzero exit code is specified, `make` prints a warning message.
`.SILENT:`	If present, `.SILENT:` has the same effect as the `-s` option.

Dependency statements

A **dependency statement** in a makefile asserts the logical relation between a target and its dependents. The syntax for a dependency statement is

target1 [*target2* ...] `:` [`:`] [*dependent1*...] [`;` *commands*] [`#` *comment*]

A sample dependency statement is

```
dancing: music.o
```

A more complex dependency statement with an associated command sequence is

```
yancy.o wally.o: defs.h ;
(TAB)      echo "defs.h has been changed"
```

A dependency statement can contain either a single colon or a double colon.

◆ **Note** A target name can appear in more than one dependency statement, but each of those statements must have the same number of colons (either one or two). ◆

Usually, dependency statements contain only a single colon. In this case, a command sequence can be associated with, at most, one dependency line; that is, a target cannot appear in more than one dependency line if there is a command sequence associated with more than one of them. For example, the fragment

```
yancy.o wally.o: defs.h
yancy.o quincy.o: menus.h
```

works because there is no command sequence associated with the dependencies in which `yancy.o` appears.

The following fragment is also correct, because there is only one command sequence associated with the dependencies in which `yancy.o` appears:

```
yancy.o wally.o: defs.h
(TAB)      echo "defs.h has been changed"
yancy.o quincy.o: menus.h
```

If the target is out of date with respect to any of the dependents on any of the lines, and a command sequence is specified (even a null one following a semicolon or tab), that command sequence is executed. Otherwise (if a command sequence is not specified), default rules can be invoked.

The following fragment uses incorrect syntax; it uses only a single colon, but a target appears in two dependency lines, each of which is associated with a command:

```
yancy.o wally.o: defs.h
(TAB)       echo "defs.h has been changed"
yancy.o quincy.o: menus.h
(TAB)       echo "menus.h has been changed"
```

In a dependency statement using two columns, a command sequence can be associated with each dependency line. For example:

```
yancy.o wally.o:: defs.h
(TAB)       echo "defs.h has been changed"
yancy.o quincy.o:: menus.h
(TAB)       echo "menus.h has been changed"
```

If the target is out of date with respect to any of the files on a particular line, `make` executes the associated commands, possibly in addition to default rules. If a target must be created, `make` executes the entire sequence of commands. This detailed form is of particular value in updating archive-type files.

If you have a single-colon and double-colon version of the same target, such as

```
x: x.o
x:: x.c
```

`make` issues a warning and continues. `make` executes the rules of the double-colon statement first and then the rules of the single-colon statement. If the double-colon rules have commands, `make` does not execute the commands associated with the single-colon rules. If this is the case, you receive a warning statement informing you that the commands are being ignored.

Commands

A **command** is usually the command line required for producing the targets from the dependents. Syntactically, a command is any string of characters, not including a number sign (`#`) (except when the `#` is in quotes) and not including a newline.

◆ **Note** When a command appears on a line separate from a dependency statement, *it must be preceded by a tab.* If not preceded by a tab, the command usually results in the message `Make: must be a separator on rules line` *x*`. Stop.` ◆

Comments

Comments are lines beginning with a number sign (#) and ending with a newline. `make` ignores these lines. `make` also ignores blank lines.

`include` lines

The C syntax for `include` lines

`#include` *include_file*

cannot be used in description files, because comments begin with a number sign. Therefore, the following policy was adopted for `include` lines in `make` description files.

If the string `include` appears as the first seven letters of a line in a makefile and is followed by a blank or a tab, `make` assumes the string following to be a filename that is to be read by the current invocation of `make`. Thus, a makefile might contain the following line:

```
include macro_defs #reads in file macro_defs
```

In this example, `macro_defs` is a file containing `make` macro definitions. No more than 16 levels of nested `include` statements are supported.

Macro definitions

Macros are defined in `make` command-line arguments or in the makefile. In the makefile, a macro definition is a line containing an equal sign, and the line must not begin with a colon or a tab. For example:

```
OBJECTS = x.o y.o z.o
```

The syntax for macro substitution is

$ (*name*)

or

$ { *name*}

The name of the macro is either a single character after the dollar sign or a *name* inside parentheses or braces. Macro names longer than one character must be put inside parentheses or braces. For example, the following macro invocations are valid:

```
$(CFLAGS)
$2
${xy}
$Z
$(Z)
```

The last two invocations listed are functionally identical. Note that two dollar signs (`$$`) can also be used to denote a dollar sign. The following fragment illustrates the assignment and use of some macros:

```
OBJECTS = x.o y.o z.o
LIBS = -lm
prog: $(OBJECTS)
(TAB)       cc $(OBJECTS) -o prog $(LIBS)
...
```

In this example, `make` loads the three object files with the math library. The command line

```
make "LIBES = -ll -lm"
```

loads them with both the `lex` (`-ll`) and the math (`-lm`) libraries.

Macro definitions on the command line override definitions in the description file, which, in turn, override the default macros.

For example, if you defined macros in your makefile, you can redefine the library on the command line for a single run of `make`, without changing the meaning of the macros defined in the makefile. For example, the command

```
make "LIBES = -lg"
```

redefines the `LIBES` macro for this run.

To see a listing of the default macros, you can consult the `Macros` part of the listing produced by the command

```
make -np
```

Internal macros

The following macros are internal and change values during the execution of a description file. These internal macros are useful generic terms for current targets and out-of-date dependents. `make` sets these internal macros as follows:

`$@`

Current target. The `$@` macro is set to the full target name of the current target. This macro is evaluated only for explicitly named dependencies. For example, in the following makefile, the current target is `zeke`, so `$@` is translated as `zeke`:

```
zeke: zeke.o
(TAB)       cc zeke.o -o $@
```

`$?`

Out of date relative to target. The `$?` macro is set to the string of names that were found to be younger than the target. This macro is evaluated when explicit rules from the makefile are evaluated. For example, the following makefile prints all files younger than `springtime`:

```
springtime:
springtime: lp $?
```

`$<`

Related file causing action. If the command was generated by a default rule, the `$<` macro expands to the name of the related dependent that caused the action. For example, the following makefile establishes an implicit rule to create targets from "`.o`" files:

```
.o:
(TAB)       cc $< -o $@
```

`$*`

Shared prefix, current, and dependent files. If the command was generated by a default rule, the `$*` macro is given the value of the filename prefix shared by the current and dependent filenames. For example, the following makefile sets the prefix `$*` to `zeke` and links `zeke.o`:

```
zeke: zeke.o
(TAB)       cc $*.o -o $*
```

In the following additions, the `D` refers to the directory part of the single-letter macro, and the `F` refers to the filename part of the single-letter macro. These are useful when building hierarchical makefiles.

`$(@D)`	current target directory
`$(@F)`	current target filename
`$(*D)`	shared directory prefix

`$(*F)`	shared filename prefix
`$(<D)`	related dependent directory
`$(<F)`	related dependent filename

For example, the following instruction uses the `D` to gain access to directory names to use the `cd` command:

```
cd $(<D); $(MAKE) $(<F)
```

Dynamic dependency parameters

The following parameters have meaning *only* within a dependency statement in a makefile.

`$$@` The current item to the left of the colon. The double dollar signs denote a metalevel macro—that is, a macro referring to another macro. Thus, `$$@` is a macro variable for whatever target is current, and `$@` is a macro for the current target. If the target is static, `$@` can be used instead of `$$@`; however, `$$@` allows for use of a dynamic target, a macro defined to denote many files, each of which is processed in turn. This is useful for building a large number of executable files, each of which has only one source file.

For example, the following makefile defines `CMDS` as the stipulated subset of single-file programs in the A/UX software command directory. Each of the programs (or `CMDS`) is compiled correctly in turn using this syntax.

```
CMDS = cat dd echo date cc cmp comm ar ld chown

$(CMDS): $$@.c
(TAB)        $(CC) -O $? -o $@
```

(See "The Default Macro Settings" for more information on `$(CC)`.)

The dependency statement for the first item in the list of `CMDS` is translated as follows:

1. The target is set to `cat`.
2. The dependent is set to `cat.c` (the current target plus `.c`).
3. The `cc` command (optimized using `-O`) runs on the dependent (`cat.c`) if it is younger than the target.
4. The results are linked into the target file (`cat`).

◆ **Note** This syntax cannot be used for multiple-file programs. To deal with multiple-file programs, you usually allocate a separate directory and write a separate makefile. Then, a specific makefile entry is made for files requiring nonstandard compilation. ◆

`$$(@F)` Another form of `$$@`, representing just the *filename* part of `$$@`. This parameter is also evaluated at execution time. For example, the following makefile maintains the `/usr/include` directory from a makefile in another directory:

```
INCDIR = /usr/include

INCLUDES = \
(TAB)       $(INCDIR)/stdio.h \
(TAB)        $(INCDIR)/pwd.h \
(TAB)        $(INCDIR)/dir.h \
(TAB)        $(INCDIR)/a.out.h

$(INCLUDES): $$(@F)
(TAB)        cp $? $@
(TAB)        chmod 0444 $@
```

The `$$(@F)` macro represents the *filename* prefix part of the current target `$@`. Because the target is also a macro, its value is equal to each of the four files named in turn. On the run of the first file,

1. The target is `stdio.h`.
2. The macro `$$(@F)` is `stdio` (the target *filename* prefix).
3. The next line copies the younger file (`$?`), if it exists, into the target file.
4. The last line changes the mode of the new target file (`$@`) (in this case, `stdio.h`) to read-only.

This pattern is repeated for the other three files stated.

Options

Suppressing printing of commands

Normally, when `make` processes a description file, each command is printed and then passed to a separate invocation of the shell after `make` substitutes for macros. The printing is suppressed in the silent mode (`make -s`), or if the special name `.SILENT` appears on a line by itself as a target in the makefile, or if the command line begins with an `@` sign. For example,

```
@size make /usr/bin/make
```

If the command line above were in a description file, the printing of the command line itself would be suppressed by the `@` sign, but the output of the command would be printed.

Ignoring errors

The `make` program normally stops if any command signals an error by returning a nonzero exit status. The `make` program ignores errors if any of the following options are used:

- The `-i` flag on the `make` command line (where the scope is global)
- The built-in target name `.IGNORE` in the description file (where the scope is the description file)
- A hyphen beginning the command string in the description file (where the scope is the command following the hyphen)

Thus, if you use the `-i` option, the target is `file.o`, and the compilation is unsuccessful, `make` effectively pretends that it worked. When `file.o` is found to be a dependent of some other files, `make` tries, for instance, to load all the object files together, and fails with an error message when one (`file.o`) is found to be missing. For all subsequent accesses (within this `make`), `make` treats `file.o` as though it existed and as though it were up to date. You should be aware of this possible consequence of the `-i` option.

Some commands return with nonzero status even though they worked correctly. For example, `diff` returns 1 to indicate the presence of differences in the compared files, and `rm` returns a nonzero status if the file you remove is already nonexistent. It is safer

to use a leading hyphen for commands that might return a nonzero exit status without indicating an error, so `make` can continue processing.

Combining commands

As stated previously, when `make` processes a description file, each command or individual command line is printed and then passed to a separate invocation of the shell after substituting for macros. Because the shell to which `make` passes each command line is a completely new invocation, you must be careful with certain commands (for example, `cd` and shell control commands) that have meaning only within a single shell process. If special means are not taken, the results of these commands are lost before the next line is executed.

One way to avoid this is to combine two or more shell commands on one line, thus keeping the same shell active on each. This can be done in one of two ways. If both commands are kept on one physical line, a semicolon (`;`) can be inserted between the commands. If the commands are put on separate physical lines but form one logical line, a semicolon (`;`) and a backslash (`\`) are supposed to be the first commands. In the latter case, the semicolon separates the commands and the backslash escapes the newline. Examples of these two methods follow:

```
# with ; both commands can be on the same line
cd ..; cc -c x.o y.o z.o

# with ; and \ before <CR>, this is read as one line

cd ..;\
  cc -c x.o y.o z.o
```

Default commands

If you need to run `make` on a file, `prog` for example, but there are no explicit commands given or relevant rules to apply, `make` looks for commands dependent on the target `.DEFAULT` to use. If there is no `.DEFAULT` target, `make` prints a message,

`Don't know how to make` *prog*`. Stop`

and stops. Thus, `.DEFAULT` can be set up by the user to specify default-case treatments for files not covered by the built-in rules of `make`. (For a listing of the types of file compilations covered by these rules, see the "Transformation Rules" section.)

Saving files

If a file or files are assigned as dependent to `.PRECIOUS`, those files are not removed, regardless of any command to the contrary. This is especially helpful to avoid the removal of targets when `make` receives an interrupt or quit.

Use of selected options

`-n` The `-n` option is useful to discover what commands `make` would execute. This option instructs `make` to print out the commands it would issue, without actually executing them.

`-t` The `-t` (`touch`) option updates the modification times on the affected files without changing anything else, and thereby can avoid a large number of superfluous recompilations. Be careful when using this option.

`-d` The `-d` (debug) option prints out a detailed description of what it is doing, including the file times. The output is verbose. Attaching a single digit to the `-d` option scales the output. If you wish to control the output, select a digit from 0 to 9. (Level 0 is minimal output, but is very clear; level 9 shows everything including flag names.)

Suffixes and rules

The `make` program uses a table of significant suffixes and a set of transformation rules to supply default dependency information and implied commands. All of this information is stored in an internal table (the default rules) that has the form of a description file. (If the `-r` option is specified, this internal table is not used.)

Suffixes

The list of suffixes is actually the dependency list for the built-in target `.SUFFIXES` in the description file. The `make` program searches for a file with any of the suffixes on the list. If such a file exists and there is a transformation rule for that combination, `make` transforms a file with one suffix into a file with another suffix.

The order of the suffix list is significant because the list is scanned from left to right. The first name formed that is associated with both a file (in the directory) and a rule (in the makefile or default rules) is made, *and no others.* The default suffix list is shown in Table 7-1.

◆ **Note** You should know the order of the default suffix list if you are not specifying a command in the makefile. Otherwise, you might make an unexpected file. ◆

Table 7-1 Default suffix list

Suffix	File type
`.o`	Object file
`.c`	C source file
`.e`	EFL source file
`.r`	`ratfor` source file
`.f`	Fortran source file
`.s`	Assembler (`as`(1)) source file
`.y`	`yacc`-C source grammar
`.yr`	`yacc-ratfor` source grammar
`.ye`	`yacc`-EFL source grammar
`.l`	`lex` source grammar

Transformation rules

`make` has an internal table of **transformation rules** that perform certain default commands if there is no command specified in the makefile. Note that the default rules also are known as the implicit rules. There are two types of transformation rules, double suffix rules and single suffix rules. In double suffix rules, `make` discerns the stage of compilation from the suffix (for example, `x.c` is a source file and `x.o` is an object file). These rules are phrased in terms of transformations from one type of suffix to another. `make` forms the names of these rules by concatenating the two filename suffixes; for example, the name of the rule to transform a `.r` file to a `.o` file is `.r.o`.

Single suffix rules describe the transformation of a file with a given suffix into one with no suffixes or a null suffix.

If a rule is listed in the internal table and there is no command sequence given in the description file, `make` uses the rule. Thus, standard transformations (from one type of file to another; for example, from a source file to an object file) do not call for a makefile entry unless nonstandard treatment is required.

If a rule is used (that is, if a default command is generated), the `$*` macro is given the value of the *filename* prefix of the file to be maintained. Then, the `$<` macro is the name of the dependent that caused the command.

`make` has all the required information for compiling programs written in languages supported by A/UX. For example, after the command

```
make x.o
```

where `x.o` is a C language object file, `make` searches for a file called `x.c` (a C language source file) in the local directory. If it finds `x.c`, `make` consults its default rules for compilation. `make` finds the rule `.c.o`, which states the default command

```
cc -O -c x.c
```

which `make` then issues to produce `x.o`.

`make` uses the default suffix list (see "Suffixes") to decide when to invoke which rules. This list tells the order in which to search for certain suffixes.

Within the `make` default rules file, the name of the rule to follow appears in the place of the target filename. Thus, the `.c.o` rule is represented by

```
.c.o:
(TAB)        cc -O -c [filename].c
```

The contents of the current default rules file used by `make` can be directed to standard output with the command

```
make -np
```

Any error messages produced at the end of this output should be ignored. The example shown in Listing 7-1 is a representative file, giving one version of the default rules used by `make`.

Listing 7-1 Sample listing of default rules file

```
# LIST OF SUFFIXES

.SUFFIXES: .o .c .c~ .y .y~ .l .l~
 .s .s~ .sh .sh~ .h .h~


# PRESET VARIABLES

 MAKE=make
 YACC=yacc
 YFLAGS=
 LEX=lex
 LFLAGS=
 LD=ld
 LDFLAGS=
 CC=cc
 CFLAGS=-o
 AS=as
 ASFLAGS=
 GET=get
 GFLAGS=


# SINGLE SUFFIX RULES

.c:
 $(CC) -n -O $< -o $@

.c~:
 $(GET) $(GFLAGS) -p $< > $*.c
 $(CC) -n -O $*.c -o $*
 -rm -f $*.c

.sh:
 cp $< $@
```

(continued)➡

```
.sh~:
 $(GET) &(GFLAGS) -p $< > .sh
 cp $* .sh $*
 -rm -f $* .sh


# DOUBLE SUFFIX RULES

.c.o:
 $(CC) $(CFLAGS) -c $<

.c~.o:
 $(GET) $(CFLAGS) -p $< > $*.c
 $(CC) $(CFLAGS) -c $*.c
 -rm -f $*.c

.c~.c:
 $(GET) $(GFLAGS) -p $< >$*.c

.s.o:
 $(AS) $(ASFLAGS) -o $@ $<

.s~.o:
 $(GET) $(GFLAGS) -p $< > $*.s
 $(AS) $(ASFLAGS) -o $*.o $*.s
 -rm -f $*.s

.y.o:
 $(YACC) $(YFLAGS) $<
 $(CC) $(CFLAGS) -c y.tab.c
 rm y.tab.o$@

.y~.o:
 $(GET) $(GFLAG) -p $< > $*.y
 $(YACC) $(YFLAGS) $*.y
 $(CC) $(CFLAG) -c y.tab.c
 rm -f y.tab $*.y
 mv y.tab.o $*.o
```

```
.l.o:
 $(LEX) $(LFLAGS) $<
 $(CC) $(CFLAGS) -c lex.yy.c
 rm lex.yy.c
 mv lex.yy.o $@

.l~.o:
 $(GET) $(GFLAGS) -p $< > $*.l
 $(LEX) $(LFLAGS) $*.l
 $(CC) $(CFLAGS) -c lex.yy.c
 rm -f lex.yy.c $*.l
 mv lex.yy.o $*.o

.y.c:
 $(YACC) $(YFLAGS) $<
 mv y.tab.c $@

.y~.c:
 $(GET) $( GFLAGS) -p $< > $*.y
 $(YACC) $(YFLAGS) $*.y
 mv -f $*.c
 -rm -f $*.y

.l.c:
 $(LEX) $<
 mv lex.yy.c$@

.c.a:
 $(CC) -c $(FLAGS) $<
 ar rv $@ $*.o
 rm -f $*.o

.c~.a:
 $(GET) $(GFLAGS) -p $< > $*.c
 $(CC) -c $(CFLAGS) $*.c
 ar rv $@ $*.o
```

(continued)➡

```
.s~.a:
 $(GET) $(GFLAGS) -p $< > $*.s
 $(AS) $(ASFLAGS) -o $*.o $*.s
 ar rv $@ $*.o
 -rm -f $*.[so]

.h~.h
 $(GET) $(GFLAGS) -p $< > $*.h
```

If two paths in the rules connecting a pair of suffixes exist, `make` uses the longer one only if the intermediate file exists or if it is named in the description file. The following examples show how this works:

1. If an `x.o` file is needed and a file called `x.c` is found in the current directory or specified in the description file, the `x.o` file is compiled using `x.c`. If an `x.l` also exists and is out of date with respect to `x.c`, that file is processed through `lex` before compiling the result. This is a case of the longer path (`x.l` to `x.c` to `x.o`) being used since the intermediate file (`x.c`) exists.
2. If the file `x.o` is needed and `x.l` but not `x.c` is found, `make` discards the intermediate C language file (in this case, `x.yy.c`) and uses the shorter path (`x.l` to `x.o`).

The default macro settings

If you know the macro names that `make` uses, you can change the names of some of the compilers used in the default rules, or the flag arguments with which they are invoked. These macro names, the default compilers they denote, and their associated flags are shown in Table 7-2.

These macros can be used as arguments on the command line to change defaults for one run of `make`. For example, the command

```
make CC=newcc ...
```

causes the `newcc` compiler to be used instead of the usual C language compiler. An example of the use of flags follows:

```
make "CFLAGS=-O" ...
```

passes the `-O` flag to the C compiler, `cc`, causing the C language optimizer to be used.

Table 7-2 Macro names and default compilers

Compiler	Macro	Flag
`make` command	MAKE	MAKEFLAGS
Assembler (`as`)	AS	-
C compiler (`cc`)	CC	CFLAGS
`ratfor` compiler	RC	RFLAGS
EFL compiler	EC	EFLAGS
`yacc`-C compiler	YACC	YFLAGS
`yacc-ratfor` compiler	YACCR	YFLAGS
`yacc`-EFL compiler	YACCE	YFLAGS
`lex` compiler	LEX	LFLAGS
`get` command	GET	GFLAGS

Sometimes it is possible to use macro redefinition instead of stating a local version of the default rule. Of course, this change is temporary because it takes place on the command line and must be restated, whenever desired, every time the file is remade. To change the `.c.o` rule you can say

```
make "CFLAGS=-V" thorax.o
```

and the option `-V` replaces the default setting for `CFLAGS` for this one run.

Changing default suffixes and rules

This section describes several ways you can modify the defaults in a makefile.

The default suffix list

You can add suffixes to the end of the default suffix list, change the order of the list, or change the contents of the list.

If you append new names to the suffix list, an entry can be included for `.SUFFIXES` in the description file. The dependents to `.SUFFIXES` are then added to the end of the default list.

To change the order or contents of the list, you must be aware that a `.SUFFIXES` line without any dependents deletes the current list of suffixes. Therefore, you must clear the current list to change the order of names. Thus, to install a new list, include lines such as

```
.SUFFIXES : # removes old list
.SUFFIXES : .n .n~ .l .l~ # installs new list
```

The default rules

You can modify or replace a default rule in a makefile. For example, if you define a `.c.o` rule in a makefile, your definition overrides the default one. For example, Listing 7-2 defines a new `.c.o` rule.

Listing 7-2 Replacing a default rule

```
.c.o: cc -V -c $< #Rule, not target
stomach.c: stomach.l #First target
stomach.l: defs.h
```

This invokes the `-V` option of `cc` every time a `.o` file is linked, printing the version of the assembler that was used.

Operation

This section describes many aspects of `make` operation, including variables, macros, precedence, and SCCS.

Environment variables

The `make` program reads environment variables from the shell and considers them in processing makefiles. These variables include `PATH`, `HOME`, `TERM`, `SHELL`, `TERMCAP`, and `LOGNAME` (see *A/UX Shells and Shell Programming* for more information on environment variables). Thus, a reference to `$(HOME)`, otherwise undefined in a makefile, is translated correctly into the full pathname for the user's home directory.

◆ **Note** The value of the `SHELL` variable determines which shell is used to execute commands in the makefile (by default, your login shell). If you want to include shell scripts that require a different shell (for example, a Bourne shell script when your login shell is the C shell), you must specify the new shell either on the command line

```
make [options] SHELL=/bin/sh
```

or you can do it by including the following line at the beginning of your description file:

```
SHELL=/bin/sh
```
◆

To see which environment variables `make` recognizes in the present directory (directed to standard output), give the command

```
make -np | head -50 | more
```

The first part of the output of this command prints the environment variables.

Macros

A macro is a variable whose value is set in a description file and can be overridden from the `make` command line. The entity that `make` terms a macro is very similar to environment variables in the shell. Although `make` and the shell use these entities in nearly identical ways, there are differences, which are described in the following paragraphs.

The following sample shell script

```
NAME=Joe
echo NAME
echo $NAME
```

produces the following result:

```
NAME
Joe
```

The difference between the first and second echo commands is that the first simply requests that the string `NAME` be echoed, while the second, through the dollar sign (`$`) requests that the contents of `NAME` be echoed. Such a request is called expansion.

Expansion is handled differently in `make`. The following example description file

```
NAME=joe
all:
        echo NAME
        echo $NAME
```

produces the following result:

```
NAME
AME
```

This is because `make` requires that macro names longer than one character be enclosed in parentheses or braces for expansion to occur. In this case, `make` sees the `$` and attempts to expand a variable named `N`. No such variable is set, so nothing is echoed and the `echo` command finishes by echoing the string `AME`. The following description file produces the desired result:

```
NAME=joe
all:
        echo NAME
        echo $(NAME)
```

The use of braces is equivalent to the use of parentheses, so that `${NAME}` is equivalent to `$(NAME)`.

Each time `make` evaluates a macro, it strips one dollar sign (`$`) from it. Therefore, an extra dollar sign should be added before any macro that is part of a shell command line. When `make` is invoked, it reads the user's environment and makes all the variables found there available for modification by the description file.

Environment variables are processed before any description file and after the built-in rules; macro definitions in a description file override environment variables of the same name. The `-e` flag option causes environment variables to override macro definitions of the same name in a description file.

The formal definition of a macro is shown here:

macro-name = *string2*

By convention, *macro-names* are uppercase. *macro-name* is an alphanumeric string that cannot contain a colon or a semicolon. The equal sign can be surrounded by spaces or tabs. *string2* is defined as all characters up to a comment character or an unescaped newline.

`make` provides several built-in macros. They include the following:

`MAKECDIR`

`MAKECDIR` is a read-only macro that expands into the full pathname of the current directory.

`MAKEFLAGS`

If not present in the environment, `make` creates the `MAKEFLAGS` macro and assigns to it the options with which `make` was invoked. `MAKEFLAGS` is processed by `make` as containing any legal input option (except `-f`, `-p`, `-P`, `-r`, and `-u`). Thus, `MAKEFLAGS` always contains the current input options. This proves very useful for large `make` commands. In fact, as noted above, when the `-n` option is used, the command `$(MAKE)` is executed anyway; hence, one can perform a make `-n` recursively on a whole software system to see what would have been executed. This is because the `-n` is put in `MAKEFLAGS` and passed to further invocations of `$(MAKE)`. This is one way of debugging all of the description files for a software project without actually causing the execution of update commands.

`MAKELEVEL`

If not present in the environment, `make` creates the `MAKELEVEL` macro, assigns an initial value of zero, and exports it. If the `MAKELEVEL` macro is already present in the environment, `make` increments its value by one. In this way, each subordinate invocation of `make` can know its level in a multilevel `make` hierarchy. This macro is read-only and cannot be modified by the description file.

`MAKEBDIR`

If not present in the environment, `make` creates the `MAKEBDIR` macro and assigns to it the absolute pathname of the current directory. If the `MAKEBDIR` macro is already present in the environment, the value is not changed. `MAKEBDIR` provides a way for each subordindate invocation of `make` to obtain the pathname of the top-level `make`.

`MAKEGOALS`

For every invocation of `make`, `make` creates the `MAKEGOALS` macro and assigns to it the targets that are specified on the command line. For the command line

```
$ make clean all clobber
```

`MAKEGOALS` is set to `clean all clobber`. If the current invocation of `make` invokes `make`, the invocation can be made as shown in the following example:

```
MAKE=make
cd dir; $(MAKE) $(MAKEGOALS)
```

In this way, the same command-line arguments can be passed to subordinate invocations of `make`.

VPATH

This version of `make` supports special processing of the macro `VPATH`, if set. `VPATH` is useful for processing files that are located in a directory other than the current directory. In the following example, `main.c` is located in the current directory. `func1.c` is located in `../common`, and `func2.c` is located in `../incl`. `make` searches the directories specified by the `VPATH` variable for any dependencies that are not in the current directory.

```
VPATH=../common:../incl
main: main.o func1.o func2.o
cc -o $@ $>
```

In this example, `$@` (described later) expands to the target name and `$>` (described later) expands to the list of dependencies on the current target. If `main.c`, `func1.c`, or `func2.c` are not present in the current directory, `make` uses its built-in rules to search for SCCS versions of the files in the current directory (see "SCCS File Handling," later). If SCCS versions of the files are not found, `make` searches the pathnames specified by `VPATH`.

The following built-in macros define values for common software generation programs or options to those programs. Description files can replace or supplement the values of these macros to change the way in which the built-in rules work:

AR	This macro is defined as `ar`.
AS	This macro is defined as `as`.
SFLAGS	This macro is defined as null and is provided as an argument to the assembler.
CC	This macro is defined as `cc`.
CFLAGS	This macro is defined as `-O` and is provided as an argument to the C compiler.
HMOD	This macro is defined as `chmod`.
CP	This macro is defined as `cp`.
F77	This macro is defined as `f77`.
F77FLAGS	This macro is defined as null and is provided as an argument to the Fortran compiler.

FORTRAN	This macro is defined as Fortran.
FORTRANFLAGS	This macro is defined as null and is provided as an argument to the Fortran compiler.
GET	This macro is defined as `get` and is used to get SCCS versions of files.
GFLAGS	This macro is defined as null and is provided as an argument to `get`.
LD	This macro is defined as `ld`.
LDFLAGS	This macro is defined as null and is provided as an argument to `ld`.
LEX	This macro is defined as `lex`.
LFLAGS	This macro is defined as null and is provided as an argument to `lex`.
MAKE	This macro is defined as `make`.
MV	This macro is defined as `mv`.
PASCAL	This macro is defined as `pascal`.
PASCALFLAGS	This macro is defined as null and is provided as an argument to `pascal`.
PC	This macro is defined as `pc`.
PCFLAGS	This macro is defined as null and is provided as an argument to `pc`.
RM	This macro is defined as `rm`.
YACC	This macro is defined as `yacc`.
YFLAGS	This macro is defined as null and is provided as an argument to `yacc`.

The following six built-in macros have special expansion capabilities that are useful for writing shell commands:

`*` The `*` macro stands for the filename part of the current dependent with the suffix deleted. It is evaluated only for built-in rules.

`@` The `@` macro stands for the full target name of the current target. It is evaluated only for explicitly named dependencies.

`<` The `<` macro is evaluated only for built-in rules or the `.DEFAULT` rule. It is the module that is out of date with respect to the target (for example, the "manufactured" dependent filename). Thus, in the `.c.o` rule, the `<` macro evaluates to the `.c` file. An example for making optimized `.o` files from `.c` files is

```
.c.o:
 cc -c -O $*.c
```

or

```
.c.o:
cc -c -O $<
```

`?` The `?` macro is evaluated when explicit rules from the description file are evaluated. It is the list of prerequisites that are out of date with respect to the target; essentially, those modules that must be rebuilt.

`%` The `%` macro is only evaluated when the target is an archive library member of the form `lib(file.o)`. In this case, `@` evaluates to `lib` and `%` evaluates to the library member, `file.o`.

`>` The `>` macro is expanded to list all the dependencies on the current rule.

These six macros can have alternative forms. When an uppercase `D` or `F` is appended to any of the six macros, the meaning is changed to "directory part" for `D` and "file part" for `F`. Thus, `$(@D)` refers to the directory part of the string `@`. If there is no directory part, `/` is generated.

The following description file demonstrates the use of the `?` and `>` macros in their standard and alternative forms:

```
pgm:
            @echo "? = $?"
            @echo "?D = $(?D)"
            @echo "?F = $(?F)"
            @echo "> = $>"
            @echo ">D = $(>D)"
            @echo ">F = $(>F)"
pgm: dir/a.o dir/b.o dir/c.o
```

When `a.o` is the only object that is newer than the `pgm`, `make` produces the following output:

```
? = dir/a.o
?D = dir
?F = a.o
> = dir/a.o dir/b.o dir/c.o
>D = dir dir dir
>F = a.o b.o c.o
```

Precedence

Each time `make` executes, `make` reads environment variables and adds them to the macro definitions. Precedence is a prime consideration in doing this properly. The following list is the default precedence of assignments:

1. command line
2. makefiles
3. environment
4. default macros

When executed, `make` assigns macro definitions in the order stated by doing the following tasks:

- *Reading the* `MAKEFLAGS` *environment variable.* Each letter in `MAKEFLAGS` is processed as an input flag argument, unless the letter is `-f`, `-p`, or `-r`. These options give directions to `make` involving overall processing, as follows:

  `-f` Precedes the makefile filename

  `-r` Leaves out the built-in rules

  `-p` Prints out all macro definitions and target descriptions

  If the `MAKEFLAGS` variable is null or is not present, `MAKEFLAGS` is set to the null string. This pass establishes the debug mode if the `-d` flag is set.

- *Reading and setting the input flags from the command line.* The command line adds to the previous settings in the `MAKEFLAGS` environment variable.

- *Reading macro definitions from the command line.* Any macro definitions set from the command line cannot be reset. Further assignments to these macro names are ignored.
- *Reading the internal list of macro definitions.* `make` reads its default rules file, which contains the internal list of macro definitions. For example, if the command

  ```
  make -r ...
  ```

  is given, and a makefile already includes all of the rules that are found in the `make` default rules file (for instance, by means of an `include` line; see "`include` Lines," earlier in this chapter), the `-r` option does not have the stated effect of "ruling out" the rules. It does not go to its default rules itself, but it cannot undo an `include` line in a makefile. In fact, the effect is identical to that occurring if both the `-r` option and the `include` line in the makefile were excluded, since they cancel each other out.
- *Reading the environment settings in the shell.* The environment variables are treated as macro definitions and marked as exported.

  ◆ **Note** Because `MAKEFLAGS` is not a variable in the `make` default rules file, this step has the effect of doing the same assignment twice. (The exception to this is when `MAKEFLAGS` is assigned on the command line.) ◆

  The `MAKEFLAGS` variable is read and set again.
- *Reading the makefiles.* Assignments in the makefiles override the environment unless the `-e` flag is used. The command line option `-e` instructs `make` to override the makefile assignments with the environment settings.

  If assigned, the `MAKEFLAGS` variable overrides the environment. This is useful for further invocations of `make` from the current makefile. There is no way to override command-line assignments. For example, if the command

  ```
  make -e ...
  ```

  is given, the variables in the environment override the definitions in the makefile and reset the precedence of assignments to the following order:

  1. command line
  2. environment
  3. makefiles
  4. default macros

This has the effect of giving the environment priority over the makefile, as opposed to the reverse in the default case.

Macro testing

`make` supports the testing of macros, where the format is:

`$(`*macro-name*`:`*test-operator*`)`

The *macro-name* can be set or unset and with or without an assigned value. The *test-operator* can be one of the following operations:

`L` The macro is expanded to the length of its contents. An empty or null value expands to zero. This test operator is useful for determining whether to examine the contents of a macro.

`V` If the macro is set and has a non-null value, the macro is expanded to null; otherwise, the macro is expanded to `#`. This test can be used to control the execution of command lines as shown here:
`$(`*macro-name*`:V)` *conditional-command*

If the macro is not set, the macro is expanded to `#`, which causes `make` to evaluate the line as a comment. As a result, *conditional-command* is not executed.

`N` If the macro is set and has a non-null value, the macro is expanded to `#`; otherwise, the macro is expanded to null. This is the opposite of the `V` test operator described earlier, although it is used in the same way as the `V` test operator.

`S` If the macro is set, the macro expands to null; otherwise, the macro is expanded to `#`.

`U` If the macro is not set, the macro expands to null; otherwise, the macro is expanded to `#`.

For example, assume you want to have a target called *clean* if the macro `$CLNFILES` is set. The dependency statement removes the files expanded from this macro. Here is how the dependency statement would look:

```
$(CLNFILES:V)clean:
            $(CLNFILES:V) @echo "Removing: $(CLNFILES)";\
            rm -f $(CLNFILES)
```

If the `$CLNFILES` macro is set and contains a non-null value, the `$(CLNFILES:V)` macro becomes null when `make` reads the description file, and the line is processed just as if the description file contained

```
clean:
            @echo "Removing: $(CLNFILES)";\
            rm -f $(CLNFILES)
```

The `$(CLNFILES)` macro is expanded just before the command line is executed. Macros that have test operators are expanded during the parsing of the command line. This means that the order of macros that have test operators is significant, which is unlike the normal behavior of macros that do not have test operators. Normal macros are expanded after all description files are read and command-line execution has begun. Expansion of macros that have test operators can be delayed by preceding additional `$` characters, just as can be done with normal macros.

In the example above, notice that `$(CLNFILES:V)` does not appear in front of each line. This is because a single command line was used, and that command line was spread over two lines, with the newline escaped by the backslash (`\`) character. If there had been multiple command lines, each command line would have to have been preceded by a `$(CLNFILES:V)` macro.

Attributes

`make` understands attributes, which can be placed before or after the dependents in a dependency list, as shown:

`target:` [*attributes*] [*dependents*] [*attributes*]

Attributes can be any of the following:

`.CURTIME` This attribute causes `make` to use the current time rather than the most recent modification time of the target, even if the target does not exist. This attribute is used with the `.FAKE` attribute to prevent the associated dependency statement from being invoked unless the dependents were updated with a newer time.

`.FAKE` If the target exists and has no dependents, the normal behavior of `make` for single-colon dependency statements is to do nothing. The addition of the `.FAKE` attribute to the dependency statement requires `make` to treat the target is if it does not exist. This, in turn, forces `make` to execute the associated commands.

`.IDEBUG`*n*: If present, `.DEBUG`*n* tells `make` not to display debugging information about this target at the desired debugging level. The variable *n* is set to a debugging level from 0 through 9. For example, to prevent this target from showing up in your debugging sections at levels 0 and 1, use `.IDEBUG0` and `.IDEBUG1`.

`.IGNORE:` This attribute causes errors from any command of the target to be ignored.

`.MAIN` The normal behavior of `make` when invoked without a target name on the command line is to search the description file for the first target, process the target, and then terminate. The addition of `.MAIN` to a dependency statement causes `make` to treat the associated dependency statement as if it were the first dependency statement in the description file.

`.PRE` This attribute informs `make` that the associated target is to be made before any others, including `.MAIN`. Hence, this attribute can be used to place initialization commands. Because the entire description file is read before the targets are processed, the placement of this attribute is position-independent within the description file.

`.POST` This attribute informs `make` that the associated target is to be processed after all others. Hence, this attribute can be used to place cleanup commands.

`.KEEPTIME` This attribute causes `make` to maintain the original modification time of the target, even after the target is regenerated.

`.OLDTIME` This attribute causes `make` to ignore the modification time of the target and apply a modification time of 0 for the purpose of determining whether the target should be updated. After the target is regenerated, `make` sets the correct modification time.

`.NOMESS` If present, `.NOMESS:` causes `make` not to echo commands or issue any warning or error messages from commands. This is useful in `PRE` and `POST` files where you might not want the user to see messages from these files.

`.NOVPATH` This attribute causes `make` to ignore the `$VPATH` macro for the associated target.

`.PRECIOUS:` With this attribute, the document is considered "precious."

`.SILENT:` With this attribute, the commands of this target are not echoed before execution.

If targets that have `.MAIN`, `.PRE`, and `.POST` attributes are dependents of other targets, the targets are made in the order dictated by the dependencies and not by the attributes.

Attributes on dependency statements with two colons apply to all of them as a unit.

Archive libraries

A `.a` suffix rule builds libraries. (There is no actual `.a` suffix appended to the filename, however; see below for how to recognize candidates for this rule.) For example, the `.c.a` rule is the rule for the following:

- compiling a C language source file
- adding a C language source file to the library
- removing the `.o` version of the C language source file

The `.y.a` rule is the rule for performing the same functions on a `yacc` file; the `.s.a` rule, for an assembler file; and the `.l.a` rule, for a `lex` file.

The current archive rules defined internally are `.c.a`, `.c~.a`, and `.s~.a`. (See the section on "SCCS Filename Prefixes" for an explanation of the tilde (~) syntax.)

Programmers might choose to define additional rules in the makefiles.

A library is then maintained with the following makefile:

```
lib: lib(ctime.o)
(TAB)       @echo lib up-to-date
```

◆ **Note** The first parenthesis in the filename identifies the target suffix rule, not an explicit `.a` suffix. ◆

For example, the actual rule `.c.a` is defined as follows:

```
.c.a:
(TAB)       $(CC) -c $(CFLAGS) $<
(TAB)       ar rv $@ $*.o
(TAB)       rm -f $*.o
```

In the `.c.a` rule:

`$@`	This macro is the `.a` target. (Using the library example, this macro is defined as `lib`.)
`$<` and `$*`	These macros are set to the out-of-date C language file, and the filename without the suffix, respectively. Using the previous example, these macros are defined as `ctime.c` and `ctime`. Using this example, the `$<` macro could have been changed to `$*.c`.

When `make` sees the `lib(ctime.o)` instruction in the makefile (assuming the object in the library is out of date with respect to `ctime.c`, and there is no `ctime.o` file), it translates that construct into the following sequence of operations:

1. `make lib`.
2. To `make lib`, `make` each dependent of `lib`.
3. `make lib(ctime.o)`.
4. To `make lib(ctime.o)`, `make` each dependent of `lib(ctime.o)`. (There are none in this example.) The following syntax allows `ctime.o` to have dependencies:

   ```
   lib(ctime.o): $(INCDIR)/stdio.h
   ```

   Thus, explicit references to `.o` files are unnecessary.

   ◆ **Note** There is also a macro for referencing the archive member name when this form is used. The `$%` macro is evaluated each time `$@` is evaluated. If there is no current archive member, `$%` is null. If an archive member exists, then `$%` evaluates to the expression between the parentheses. ◆

5. Use default rules to try to build `lib(ctime.o)`. (There is no explicit rule.)

   ◆ **Note** It is the first parenthesis in the name `lib(ctime.o)` that identifies the (`.a`) target suffix. This is the key. There is no explicit `.a` at the end of the `lib` library name. The parenthesis forces the `.a` suffix. In this sense, the suffix is hard-wired into `make`. ◆

6. Break the name `lib(ctime.o)` into `lib` and `ctime.o`. Define two macros, `$@` (`=lib`) and `$*` (`=ctime`).

7. Look for a rule `.X.a` and a file `$*.X`. The first `.X` (in the `.SUFFIXES` list in the default rules file) that fulfills these conditions is `.c`, so the rule is `.c.a` and the file is `ctime.c`.
8. Set `$<` to `ctime.c` and execute the rule. In fact, `make` must then make `ctime.c`. The search of the current directory yields no other candidates, however, and the search ends.
9. The library is updated. Perform the next instruction associated with the `lib:` dependency. Therefore, `make` echos

   ```
   lib up-to-date
   ```

SCCS files

`make` can be used on SCCS files and can run `get` on them, if required, before otherwise processing them. Those unfamiliar with SCCS (Source Code Control System) should refer to Chapter 8, "SCCS Reference."

SCCS filename prefixes

`make` syntax does not allow for direct prefix references except with SCCS files.

SCCS filenames are preceded by an `s.` prefix. `make` uses a tilde (~) appended to the suffix to identify SCCS files. The expression `.c~.o` refers to the rule that transforms an SCCS C language source file into an object file.

The following example shows a transformation from an SCCS filename to a name with a suffix already fixed for `make`: the SCCS filename `s.file1.c` into the non-SCCS, `make`-ready filename `file1.c~`. This file is then assembled using the command

```
.c~.o:
(TAB)       $(GET) $(GFLAGS) -p $< > $*.c
(TAB)       $(CC) $(CFLAGS) -c $*.c
(TAB)       -rm -f $*.c
```

The tilde appended to any suffix transforms the file search into a search for an SCCS filename with the actual suffix named by the dot and all characters up to (but not including) the tilde (~).

SCCS filename suffixes

The following SCCS suffixes are internally defined:

```
.c~ .y~ .s~ .sh~ .h ~
```

SCCS transformation rules

The following rules involving SCCS transformations are internally defined:

```
.c~: .l~.o: .sh~: .y~.c: .c~.o:

.c~.a: .s~.o: .s~.a: .y~.o: .h~.h:
```

These rules transform SCCS files into non-SCCS format and perform the compilations indicated by the letter combinations in the rule names. (See "Transformation Rules" for how to translate rule names into the rules they designate.)

Other rules and suffixes that might prove useful can be defined using the tilde as a handle on the SCCS filename format.

SCCS makefiles

SCCS makefiles are "invisible" to `make` in that if you give the command

```
make
```

and only a makefile named `s.makefile` resides in the current directory, `make` will `get`, read, and remove the file. `get` creates a file called `makefile` that remains in the directory (in addition to the *p-file*, `p.makefile`). If the `-f` option is used, `make` will `get`, read, and remove the specified makefile (as well as `include` files), creating a non-SCCS makefile named the same as the old SCCS version, except that the `s.` prefix is removed.

Advanced topics

This section details additional capabilities of `make`. The topics include maneuvering through directories, using shell scripts with `make`, and dynamic `include` file dependency generation.

Walking the directory tree

It is possible to get `make` to walk the directory tree, either by guiding it explicitly or by including a shell script that discovers, implicitly, what directories exist, so that it can visit them. While `make` is in each directory, it can `make` the files specified in the directory makefile. This allows you to bring whole systems up to date by having `make` follow directions in one local (meta-)makefile instead of you having to change directories yourself.

The explicit route is, by far, the easiest. If you know the structure of your tree and the names of all the directories you need to use, you can include commands in a makefile in the directory at the top of your tree. If, below your current directory, you have directories named `io`, `os`, and others, you can include lines like the following ones in your makefile:

```
all:
(TAB)       cd io; make -f io.mk; \
(TAB)       cd ../os; make -f os.mk;
```

The backslash (`\`) at the end of command lines is necessary if you want to keep the same invocation of the shell active for a group of commands. If a different shell is invoked, the directory information is lost.

If, for example, no backslash terminates the first command line, and so a different shell was invoked on the second line, the second `cd` would be executed from the parent directory for `io` and `os` instead of from the `io` directory. In this case, to keep the same effect, the line should read

```
(TAB)       cd os; make -f os.mk;
```

As this shows, it is possible to write a script that does invoke a new shell with each line and still travels the directory tree. This just changes the mode of travel: With the one-shell-per-journey method, you state explicit directions for going to each directory from where you are relative to that directory *and for going back to the originating directory afterward.* With the one-shell-per-command method, you state explicit directions (that is,

a full pathname) for going to the directory, and the return trip is done for you when the shell you are using quits.

To travel a tree of unknown structure but with fairly standard makefile names (like *dirname*.`mk`, where *dirname* stands for the name of the directory where the file is located), you could use a fragment like the following one in your makefile:

```
subdirs:
(TAB)       for i in 'find /pathname -type d -print'; \
(TAB)       do \
(TAB)         if test -f $$i/$$i.mk; \
(TAB)         then \
(TAB)               cd $$i; \
(TAB)               $(MAKE) -f $$i.mk; \
(TAB)         fi \
(TAB)       done
```

◆ **Note** The code section above is a Bourne shell script, and it works only if your login shell is `/bin/sh` or your `SHELL` environment variable is set to `/bin/sh`. See "Environment Variables" for more information on using different shells to execute a makefile. ◆

The `make` predecessor tree

The `$!` macro represents the current predecessor tree. A `make` predecessor tree contains the series of files linked through the dependency relation for one run of `make`. For example, using the `makefile`

```
all: cat
(TAB)       @echo cat up-to-date
cat: cat.c
(TAB)       echo $!
```

when the command `echo $!` is executed, the variable `$!` evaluates to

```
cat.c cat all
```

which is the current predecessor tree of this run of `make`, read from left to right (leaf to root, respectively). The connection constituting branches is the "is depended on by" relation: The left-most file is depended on by the next file to the right, and so on. Thus, the nodes are dependents of their right neighbors and are targets of their left neighbors (except for the leaf). The predecessor tree can be useful as a debugging tool for `make` itself, if what it has done does not make sense. Examination of the tree can reveal why certain files were updated, or which files were touched in this run of `make`.

Another means of debugging must be found if `make` prints the following message:

```
$! nulled, predecessor circle
```

If the predecessors of a file are circular, they cannot form a tree, and one will not be printed. The actual evaluation of the `$!` macro is terminated, and the macro value is set to null.

The makefile as shell script

If a target cannot be found in the local or specified directory, `make` attempts to create it. When `make` discovers the absence of the file corresponding to *target*, it considers *target* to be out of date and so executes the specified command sequence. If the results do not include creating the target, this leaves the directory in question in the same state, ready for the same scenario to take place whenever the `make` command is invoked.

This allows a makefile to function more like a shell script, with each absent target causing `make` to try to create it, using the command sequence specified.

Unintended targets

`make` considers missing files to be out of date and processes them. Conversely, existing files might be mistakenly deemed up to date (because of user error) and skipped for processing by `make`. This might happen in the situation described in "The Makefile as Shell Script" if one of the targets was

```
print:
(TAB)       lp foo bazz fizz
```

Here, the command sequence creates no file called `print`, so the same description file can be used over and over for maintenance, each time executing this line. If, however, you inadvertently name a program in that directory `print`, this latter file's modification information is checked to determine whether `print` needs to be remade. `make` will probably find `print` to be up to date, and tell you so on the screen. Failure to note this might cause a bug that is hard to trace in the working of the shell script description file, even though the entry for `print` is correct.

Mnemonic targets

A useful method of using `make` is to include targets with mnemonic names and commands that do not actually produce a file with the same name as the label in the shell script. These entries can take advantage of the ability in `make` to generate files and substitute macros. For example, `save` might be included to copy a certain set of files, or an entry `cleanup` might be used to throw away unneeded intermediate files. It is also possible to maintain a zero-length file purely to keep track of the time at which certain commands were performed. For example,

```
print: $(FILES)
(TAB)       pr $? | lp
(TAB)       touch print
```

The `print` entry prints only the files changed since the last `make print` command. A zero-length file `print` is maintained to keep track of the time of the printing, the time since the file `print` was last touched. The `$?` macro in the command sequence then picks up only the names of those files changed since `print` was touched. The `touch` command creates this zero-length file if no file called `print` exists in this directory.

Macro translation

To supplement macro definition and substitution, `make` also provides a macro translation facility. As a macro is evaluated, the translation takes place within the set of names of items to which the macro refers. (Such item names are probably filenames; in any case, they are considered as strings, where a string is delimited by blanks or tabs.) Thus, the macro translation facility allows for more refined and narrow macro definitions and for more concise code in description file command sequences.

The formats for macro translation follow:

$(*macro-name*:*string1*=*string2*)

This tells `make` to substitute *string2* for *string1* everywhere among the item names produced on evaluation of *macro-name*. The `make` utility attempts to assume that these substitution strings are suffixes; however, a substitute sequence can be any number of the trailing characters of *string1*. For example:

```
SAMPLE=/a/b/file.test

all:
      @echo "1 $(SAMPLE:file=FILE)"
      @echo "2 $(SAMPLE:test=TEST)"
      @echo "3 $(SAMPLE:a/=A/)"
      @echo "4 $(SAMPLE:b/file.test=K)"
      @echo "5 $(SAMPLE:a=A)
```

has the following output

```
1 /a/b/file.test
2 /a/b/file.TEST
3 /a/b/file.test
4 /a/K
5 /a/b/file.test
```

In the preceding example, only the second and fourth examples are successful. The other examples fail because they do not substitute the trailing characters of the expanded macro.

The following example demonstrates the usefulness of string substitution:

```
all: /u/test/a.o
      cc -S $(?:a=.c)
      mv $(?:.o=.s) .tmp
      sed "s/text/data/" > $(?:.o=.s) < .tmp
      as -o $@ $(?:.o=.s)
      rm .tmp
```

The preceding example uses the `@` (expand to the full target name of the current target) and the `?` (expand to the list of out-of-date dependencies) macros to produce the assembly language file for each dependent of the `all` target, change each occurrence of text to data using `sed`, and assemble each resulting `.s` file.

Substitution can even work on macros that are part of shell command lines. This version of `make` supports substitutions of macros that are part of dependency lists.

Another form of macro expansion similar in style to `ed`(1) substitution is

`$(`*name*`:/`*regular_expression*`/`*replacement_text*`/)`

First, the *name* variable is expanded and every occurrence of *regular_expression* is substituted with the *replacement_text*. For example, if you have the following makefile:

```
VAR = %file1 %file2
all:
            @echo "test 1: $(VAR:/%//)"
            @echo "test 2: $(VAR:/%/&/)"
            @echo "test 3: $(VAR:/%/X&/)"
```

running the `make` command produces

```
test 1: file1 file2
test 2: %file1 %file2
test 3: X%file1 X%file2
```

Another form of macro expansion involves pattern-matching. The expansion is in the form

`$(`*variable*`:=`*pattern*`)`

The *variable* is expanded into words separated by white space and all of the words that do not match the pattern are removed. Using this macro translation on the makefile

```
WORDS= One Two Three Four Five
all:
      @echo "Words with 'o' in them: ${WORDS:=*o*}"
```

produces this output

```
Words with 'o' in them: Two Four
```

The final form of macro expansion is

$(*name*::*default*)

The *name* variable is expanded and if this variable is undefined or NULL, the *default* value is returned. Otherwise the value of *name* is used. For example, the makefile

```
FLG=${CFLAGS::-O}
all:
            @echo "FLG: ${FLG}"
```

produces the following output when just the `make` command is used

```
FLG: -O
```

and produces the following output when the `make` command is used with the argument `CFLAGS=`

```
# make CFLAGS=
FLG: -O
```

and produces the following when the `make` command is used with the argument `CFLAGS=-F`

```
# make CFLAGS=-F
FLG: -F
```

Dynamic Include File Dependency Generation

The `make` utility includes the ability to examine selected source files for the `#include` directives. These include files are added to the target dependency list. This feature relieves you from having to set up and create the include file dependency list.

The only disadvantage to having `make` create the dependency list is that some include files might be placed on the target dependency list that would normally be left out during compilation because of an `#ifdef`. However, this does not cause any problems; the target is still updated properly. The `make` utility does not add an include file to the target dependency list unless that include file really exists, so no damage can result.

The Dynamic Include File Dependency Generation (DIFDG) is enabled by defining the `_MAKE_DIFDG_SUFFIXES` variable with a list of source file suffixes to be searched, as in this example, or by use of the `-G` (generate) option to `make`:

```
_MAKE_DIFDG_SUFFIXES= .c .s .f .p .l .y
```

The _MAKE_DIFDG_SUFFIXES variable must contain at least one suffix to enable DIFDG. An empty variable here does not have added meaning. The suffixes that are ignored are .o, .h, and .a.

The list of directories to search for these include files can be specified with the _MAKE_DIFDG_INCDIRS variable. The order is important because make searches each directory for include files until the files are found. Just like cpp(1), make looks for include files of the form header.h first in the same directory as the source file (not always the current directory), then in the directories listed in the _MAKE_DIFDG_SUFFIXES variable. If the include file has the form <header.h>, the only directories searched are those listed in the variable. If this variable is defined but not assigned a value, the only directory that is searched is the source file directory. This means <header.h> forms always fail because there is not a directory to search. If this variable isn't defined, make uses the default /usr/include directory.

A prefix can be added to each include file dependent whose full pathname starts with /usr/include, by use of the _MAKE_DIFDG_PREFIX variable. This is only used when the user requested that the include file dependencies be written using the -M (map) option. There is no default. An example is

```
_MAKE_DIFDG_PREFIX= $$(SGS_INCDIR)
```

An alternate way of creating include file dependency files is with the C preprocessor, cpp. This method is much slower than make. The last variable for DIFDG, _MAKE_DIFDG_CPPFLAGS, is defined with the flags to be passed to cpp. The mere defining of this variable enables the cpp method of finding include files. Otherwise, the faster version is used. When you assign a value to this variable, keep in mind that only words that start with a hyphen are passed to cpp, as in this example (it is assumed that DEFINES is a variable that contains -D style macros):

```
_MAKE_DIFDG_CPPFLAGS= -Y $(DEFINES)
```

If the _MAKE_DIFDG_FILE variable is set and non-null, and DIFDG is enabled, the DIFDG include file dependencies are written to it when make exits.

If the -G option is used, the defaults are

```
_MAKE_DIFDG_SUFFIXES= .c .l .y
_MAKE_DIFDG_INCDIRS= /usr/include
_MAKE_DIFDG_CPPFLAGS=
_MAKE_DIFDG_PREFIX=
```

A warning for system administrators

If the system setting for `date` is wrong (especially if it is very far behind the actual date), `make` issues a warning message. (Dates are automatically considered incorrect if they are before 1970.) Since `make` works by comparing previous dates with the current one, it is important to make sure that the current system date is accurate. To ensure proper functioning of `make`, the accuracy of `date` should be checked frequently. Also check the accuracy of the system clock in the Control Panel. If necessary, reset this clock as well, to reflect the correct date.

8 SCCS Reference

SCCS for beginners / 8-3
SCCS files / 8-7
SCCS command conventions / 8-16
SCCS command summary / 8-22

The Source Code Control System (SCCS) is a collection of A/UX commands that controls and reports on changes to files of text. SCCS is a valuable tool for version management of program source code or ordinary text files. In large group projects, SCCS prevents multiple, inconsistent versions of files from accumulating in several places. For a single user, multiple versions of a file can be stored without using a lot of disk space, previous versions can be easily reconstructed, and versions can be tracked with a simple, consistent numbering scheme. SCCS provides facilities for

- efficient storage of multiple versions of files
- retrieving earlier versions of files
- controlling update privileges to files
- identifying the version of a retrieved file
- recording when, where, why, and by whom each change was made to a file

SCCS stores the original file on disk. Whenever changes are made to the file, SCCS stores only the changes. Each set of changes is called a **delta.** When you retrieve a particular version of the file (the default is the most recent version), SCCS applies the appropriate deltas to the original file to reconstruct that version.

This chapter provides an introduction and a general reference guide to SCCS. The following topics are covered here:

- *SCCS for beginners* A step-by-step guide to creating SCCS files, updating them, and retrieving a version of a file.
- *SCCS files* A description of the protection mechanisms, format, auditing, and delta numbering of SCCS files. The differences between individual SCCS use and group or project SCCS use are discussed, and the role of the SCCS administrator in a group project is introduced.
- *SCCS command conventions* A description of the conventions that generally apply to SCCS commands and the temporary files created by SCCS commands.
- *SCCS command summary* A summary of SCCS commands and their arguments.

In addition to the programs described in this chapter, the `sccs` command provides a front end to SCCS functionality. Basically, the `sccs` front end runs the SCCS commands documented in the "SCCS Command Summary" as well as several commands that are equivalent but easier to use than the most frequently used SCCS commands. See `sccs`(1) in *A/UX Command Reference* for more information on the `sccs` front end.

Creating an SCCS file

Using a text editor, create an ordinary text file named `lang` that contains a list of some programming languages:

```
C
PL/I
FORTRAN
COBOL
ALGOL
```

To bring the tools of SCCS into play, you need to create a (different) file that various SCCS commands can read and modify. You can do this with the `admin` command, as follows:

```
admin -ilang s.lang
```

The `admin` command with the `-i` keyletter (and its value, `lang`) creates a new SCCS file and initializes it with the contents of the file named `lang`. An initial SCCS delta is created by applying a set of changes (the contents of `lang`) to a new (null) SCCS file (`s.lang`).

All SCCS files must have names that begin with "`s.`". This effectively limits SCCS filenames to 12 characters.

Each delta is assigned a name called the SCCS identification string, or **SID.** The SID is normally composed of two components (the release number and the level number) separated by a period. For example, the initial version of a file is delta 1.1 (that is, release 1, level 1). SCCS keeps track of subsequent versions of a file by incrementing the level number whenever you create a new delta. The release number can also be changed (allowing, for example, deltas 2.1, 3.1, and so on) to indicate a major change to the file.

The `admin` command returns a warning message (which also can be issued by other SCCS commands):

```
No id keywords (cm7)
```

The absence of keywords is not a fatal error under most conditions, and this warning message does not affect the SCCS file you have created. In the following examples, this warning message is not shown, although it might actually be issued by the commands.

You should now remove the `lang` file from your directory:

```
rm lang
```

Retrieving a file and storing a new version

To reconstruct the `lang` file you just deleted, use the SCCS get command:

```
get s.lang
```

This retrieves the most recent version of file `s.lang` and prints the messages

```
1.1
5 lines
```

(the SID of the version retrieved, and the length of the retrieved text). The retrieved text is placed in another file called the *g-file*. The name of the *g-file* is formed by deleting the `s.` prefix from the name of the SCCS file. Hence, the file `lang` is reconstructed.

When you use the `get` command with no keyletters (in the format above) the `lang` file is created with read-only mode (mode 440), and no information about the SCCS file is retained. If you want to be able to change an SCCS file and create a new version, use the `-e` (edit) keyletter on the `get` command line:

```
get -e s.lang
```

The `-e` keyletter causes `get` to create `lang` with read-write permission and places certain information about the SCCS file in another file called the *p-file*, which is read by the `delta` command when the time comes to create a new delta. The same messages are displayed, as well as the SID of the next delta (to be created). For example:

```
get -e s.lang
```

produces

```
1.1
new delta 1.2
5 lines
```

After this command, you can edit the `lang` file and make changes. For example, suppose that you use `vi` to create the following new version of the file:

```
C
PL/I
FORTRAN
COBOL
ALGOL
ADA
PASCAL
```

The command

```
delta s.lang
```

records the changes you made to the `lang` file within the SCCS file. SCCS prints the message

```
comments?
```

Your response should be a description of why the changes were made. For example:

```
comments? added more languages
```

The `delta` command then reads the *p-file* and determines what changes were made to the file `lang`. When this process is complete, the changes to `lang` are stored in `s.lang`, and `delta` displays

```
1.2
2 inserted
0 deleted
5 unchanged
```

The number 1.2 is the SID of the new delta, and the next three lines refer to the changes recorded in the `s.lang` file.

Retrieving versions

The `-r` keyletter allows you to retrieve a particular delta by specifying its SID on the `get` command line. For the previous example, the following commands are all equivalent:

```
get s.lang
get -r1 s.lang
get -r1.2 s.lang
```

The numbers following the `-r` keyletter are SIDs.

The first command retrieves the most recent version of the SCCS file, because no SID is specified. When you omit the level number of the SID (as in the second command), SCCS retrieves the most recent level number in that release (in the previous example, the latest version in release 1, namely 1.2). The third command explicitly requests the retrieval of a particular version (in this case, also 1.2).

Whenever a major change is made to a file, the significance of that change is usually indicated by changing the release number (the first component of the SID) of the delta being made. Because normal automatic numbering of deltas proceeds by incrementing the level number (the second component of the SID), you must explicitly change the release number as follows:

```
get -e -r2 s.lang
```

Because release 2 does not yet exist, `get` retrieves the latest version *before* release 2 and changes the release number of the next delta to 2, naming it 2.1 rather than 1.3. This information is stored in the *p-file* so the next execution of the `delta` command produces a delta with the new release number. The `get` command then produces

```
1.2
new delta 2.1
7 lines
```

which indicates that version 1.2 is retrieved and that 2.1 is the version `delta` creates. Subsequent versions of the file are created in release 2 (deltas 2.2, 2.3, and so on).

On-line information

The `help` command is useful whenever there is any doubt about the meaning of an SCCS message. Detailed explanations of almost all SCCS messages can be found using the `help` command and the code printed in parentheses after the message.

If you give the command

```
get abc
```

SCCS prints the message

```
ERROR [abc]: not an SCCS file (co1)
```

The string `co1` is a code that can be used to obtain a fuller explanation of that message using the `help` command. The command

```
help co1
```

produces

```
co1:
"not an SCCS file"
A file that you think is an SCCS file
does not begin with the characters "s.".
```

SCCS files

This section discusses the protection mechanisms used by SCCS, the format of SCCS files, recommended procedures for auditing SCCS files, and how deltas are numbered.

Standard A/UX protection

In addition to the special SCCS flags and keyletters described in the next section, "SCCS Protection Mechanisms," SCCS uses standard A/UX protection mechanisms to prevent you from making changes to SCCS files using non-SCCS commands. The following precautions are automatically taken by SCCS:

- When you create an SCCS file (using `admin`), it is automatically given mode 444 (read-only) if your `umask` is less than or equal to 333. If your `umask` is 334, the SCCS file is created with mode 440 (no read permission for others). If your `umask` is 344, the SCCS file is created with mode 400 (read permission for the owner only). If your `umask` is 444 or higher, the SCCS file is created with no permissions across the board, and a lock file, also called a *z-file*, is created. The preferred mode for an SCCS file is 444; this protects against modifying SCCS files using non-SCCS commands and should not be changed.
- If you make a hard link from an SCCS file to another file, SCCS commands do not process the SCCS file. SCCS commands produce an error message rather than process a file that has been linked. The reason for this is the same: Protection is provided against using non-SCCS commands to modify SCCS files.

SCCS protection mechanisms

SCCS provides the following protection features directly: three SCCS file flags (release ceiling, release floor, and release lock) and a user list for SCCS files.

The SCCS file flags are set using the `-f` keyletter with the `admin` command. This keyletter specifies a flag and possibly a value for the flag, to be placed in the SCCS file. Several `-f` keyletters can be supplied on a single `admin` command line (see "SCCS Flags" under "Create SCCS Files: `admin`" later in this chapter).

The flags used for file protection are

`c` *ceiling* — The highest release *(ceiling)* that can be retrieved by a `get` command for editing. *ceiling* is a number less than or equal to 9999. If this flag is not used, the default value for *ceiling* is 9999, which allows all releases up to and including 9999 to be retrieved for editing.

`f` *floor* — The lowest release *(floor)* that can be retrieved by a `get` command for editing. *floor* is a number less than 9999 and greater than 0. If this flag is not used, the default value for *floor* is 1, which allows the first release to be retrieved for editing.

`l` *list* — A list of locked releases to which deltas can no longer be made. (See `admin`(1) in *A/UX Command Reference* for the complete syntax of this list.) The `get -e` command fails if you attempt to retrieve one of these locked releases for editing. The character `a` in *list* can be specified to protect all releases for the named SCCS file.

SCCS files can also contain a user list of login names and/or group IDs of users who are or are not allowed to create deltas of that file. This list is empty by default, which means that anyone can create deltas. To add names to the list (either to allow permission or to deny it) the `-a` keyletter is used with the `admin` command. The argument to the `-a` keyletter can be

login-name — A login name or numerical group ID can be specified; a group ID is equivalent to specifying all login names common to that ID.

`!` *login-name* — If a login or group ID is preceded by an exclamation character (`!`), it is denied permission to make deltas.

These features are described in more detail under the `admin` command.

Administering SCCS

If you are using SCCS to manage personal files, the protection mechanisms described in the previous section should be used to keep certain releases from being modified, or to prevent you from accidentally modifying your files without using SCCS.

Aside from these protections, you can simply use SCCS directly. See "Delta Numbering" later in this chapter for information on storing and retrieving different releases.

Group project administration

If you are using SCCS to manage and protect files in a large project with several users having access to the same files, a single user should own the SCCS files and directories. This single user is the only one to administer the SCCS files.

The following precautions are recommended:

- Directories containing SCCS files should be mode 755. This allows only the owner of the directory to modify its contents.
- SCCS files should be kept in directories that contain only SCCS files (and any temporary files created by SCCS commands). This simplifies protection and auditing of SCCS files. The contents of directories should correspond to convenient logical groupings—for example, subsystems of a large project.
- No SCCS users other than the SCCS administrator should be able to use those commands that require write permission in the directory containing the SCCS files. Instead, a project-dependent program should be written to provide an interface to certain SCCS commands, usually the `get`, `delta`, and, if desired, `rmdel` and `cdc` commands.

This last precaution requires that you write an interface program (usually specific to the project) that invokes the desired SCCS command and gives other users (who are not the owners of the SCCS files) the permissions they need to modify specific SCCS files, using only those commands that are linked to the interface program.

◆ **Note** If you are not using the `sccs` front end (see `sccs`(1) in *A/UX Command Reference*), you might need to write an interface program such as the sample program shown in Listing 8-1 to handle special file permissions for a particular project. ◆

The sample program in Listing 8-1 causes the invoked command to inherit the privileges of the interface program for the duration of the execution of that command. Users whose login names or group IDs are in the user list for that file (but who are not the owner), and who have the path to the executable interface program in their `PATH` variable, are given the necessary permissions only for the duration of the execution of the interface program. They can modify the SCCS files only through the use of those commands that are linked to the interface program.

Listing 8-1 Sample interface program for group projects

```
main (argc, argv)
int argc;
char *argv[];
{
   register int i;
   char cmdstr [BUFSIZ];

   /* Process file arguments
      (those that don't begin with '-') */
   for (i = 1; i < argc; i ++)
      if (argv [i][0] != '-')
         argv[i] = filearg (argv[i]);

   /* Get 'simple name' of name
      used to invoke program
      (strip off directory prefix, if any) */
   argv[0] = sname (argv[0]);

   /* Invoke actual SCCS command,
      passing arguments */
   sprintf(cmdstr, "/usr/bin/%s", argv[0]);
   execv(cmdstr,argv);
}
```

This sample interface program is an example only; the functions `sname` and `filearg` are not standard functions. You should write these and any other functions required by your project.

Such an interface program must be owned by the SCCS administrator, must be executable by the new owner, and must have the `setuid` (set user ID on execution) bit on (see `setuid`(2)).

Links can then be created between the executable interface program and the command names. For example, if the path to the file is

```
/sccs/interface.c
```

then the commands

```
cd /sccs
cc interface.c -o inter
```

compile the program into the executable module `inter`. At this point, the command

```
chmod 4755 inter
```

sets the proper mode and `setuid` bit. You can then create links from any directory with the commands

```
ln /sccs/inter get
ln /sccs/inter delta
ln /sccs/inter rmdel
ln /sccs/inter cdc
```

The full pathname of the directory containing the links must then be included prior to the `/usr/bin` directory in the `PATH` variable (in the `.profile` or `.login` files of all SCCS users who need to use the desired SCCS commands). For example,

```
PATH=(.:/usr/new:/bin:/sccs:/usr/bin)
```

Depending on the type of interface program you wrote, the names of the links can be arbitrary (if the program can determine from them the names of the commands to be invoked), the pathname to your project can be supplied, and so on. If the pathname to your project is supplied in the interface program, the user can use the syntax

```
get -e s.abc
```

regardless of where the user is currently located in the file system.

SCCS file formats

SCCS files are composed of ASCII text arranged in six parts, as follows:

checksum	This part of the file contains the sum of the ASCII values of all characters in the file (not including the checksum itself). The SCCS checksum is described in "SCCS File Auditing."
delta table	This part contains information about each delta, such as type, SID, date and time of creation, and commentary.
user list	This is a list of login names and/or group IDs of users who are allowed to modify the file by adding or removing deltas. The user list is described under "SCCS Protection Mechanisms."
flags	This part contains indicators that control certain actions of SCCS commands. Flags are discussed under "Create SCCS Files: `admin`."
descriptive text	This is arbitrary text provided by the user, usually comments that provide a summary of the contents and purpose of the file. Descriptive text is discussed under "Create SCCS Files: `admin`."
body	This is the actual text of the ASCII file being administered by SCCS, intermixed with internal SCCS control lines.

For information regarding the physical layout of SCCS files, see `sccsfile`(4) in *A/UX Command Reference.*

◆ **Note** Because SCCS files are ASCII files, they can be processed by other A/UX commands such as `vi`, `grep`, and `cat`. This can be convenient when an SCCS file must be modified manually or when you simply want to look at the file. However, it is extremely important to be careful about introducing changes that affect future deltas. It is wise to make a backup copy first. ◆

SCCS file auditing

On rare occasions (such as a system crash), an SCCS file might be destroyed or corrupted (that is, one or more blocks of it might be destroyed). If the entire SCCS file has been trashed, the SCCS commands issue an error message when you attempt to process that file. In this case, you need to restore the file from your most recent backup copy.

If one or more blocks of an SCCS file are trashed by a system crash, the SCCS commands recognize this through an inconsistent checksum. In this case, the only SCCS command that processes the file is the `admin` command with the `-h` or `-z` keyletter:

`admin -h s.`*file1* `s.`*file2* ...

It is a good idea to use these commands routinely to audit your SCCS files to detect any inconsistent checksums (indicating file corruptions). If the new checksum of any file is not equal to the checksum in the first line of that file, SCCS prints the message

`corrupted file (co6)`

This process continues until all the files are examined. The `admin -h` command also can be applied to directories:

`admin -h` *directory1 directory2* ...

This prints an error message for any corrupted files, but does not automatically report missing SCCS files. To determine whether any of your SCCS files are missing, list the contents of each directory (`ls`).

If you have an SCCS file that is extensively corrupted, the best solution is to restore the file from your most recent backup copy. If there is only minor damage, you might be able to repair it using a text editor. In this case, after you repair the file, use the command

`admin -z s.`*file*

This recomputes the checksum of the file so that it agrees with the file contents. After you use `admin -z`, any corruption that existed in the file is no longer detectable by the `admin -h` command.

Delta numbering

SCCS deltas are changes applied to an original (null) file to produce different versions and releases of your file.

SCCS names deltas with an SCCS identification string (a SID). SIDs have exactly two components (the *release* number and the *level* number) separated by a period:

release`.`*level*

SCCS names the initial delta 1.1. This is considered a set of changes applied to the null file. Subsequent deltas are named by incrementing the level number (1.2, 1.3, and so on) when the delta is created. If you make a major change to the file, you might want to

specify a new release number when you create the new delta. In this case, SCCS assigns a new release number (2.1) and subsequent deltas are incremented as in release 1. This is shown in Figure 8-1.

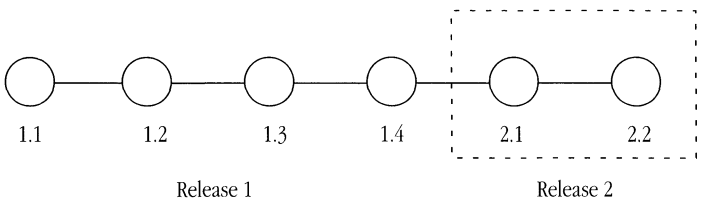

Figure 8-1 A linear progression of versions

In this simplest case, the deltas progress linearly; that is, any delta is dependent on all preceding deltas. When SCCS reconstructs a particular version of your SCCS file, it applies all deltas up to and including the number you specify. In most cases, this is all you need to know about SCCS delta numbering.

Branch deltas

The linear progression of file versions shown above is sometimes called the *trunk* of the SCCS tree for that file. Under special conditions, you may need to use a *branch* in the tree: an independent progression of deltas that does *not* depend on all previous deltas for that file.

For example, suppose you have a program at version 1.3 that is being used in a production environment. You are developing a new release (release 2) of the program, and already have several deltas of that release. This situation uses the simple linear organization shown above.

Now suppose that a user reports a problem in version 1.3 which requires changes only to version 1.3 but does not affect subsequent deltas. This requires a branch from the previous linear ordering. The new (branch) delta name consists of exactly four components: release and level numbers (as in the trunk delta) plus a branch number and sequence number, all separated by periods:

release. *level*. *branch*. *sequence*

Thus, a branch delta can always be identified as such from its name.

Once you create a branch delta, SCCS increments subsequent deltas on that branch by incrementing the sequence number. This is shown in Figure 8-2.

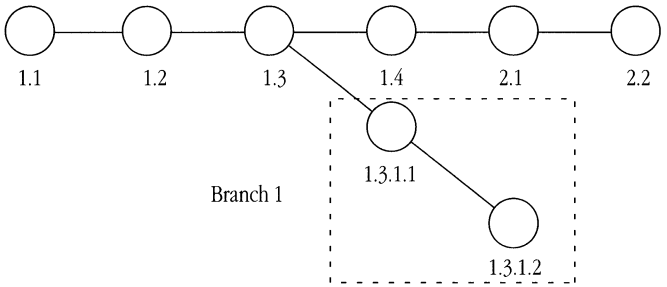

Figure 8-2 A branching SCCS tree

While SCCS increments the sequence number on each branch, it increments the branch number according to *when you create the branch.* If you need to complicate your SCCS branch structure, consider this carefully. While the trunk delta (the initial linear progression) can always be identified by the branch delta name (by the release and level numbers), it is not possible to determine the entire path leading from the trunk delta to the particular branch delta you might have retrieved.

For example, if delta 1.3 has one branch, all deltas on that branch are named 1.3.1.*n*. If a delta on this branch (for example, delta 1.3.1.1) has a branch, all deltas on the new branch are named 1.3.2.*n*. This is shown in Figure 8-3.

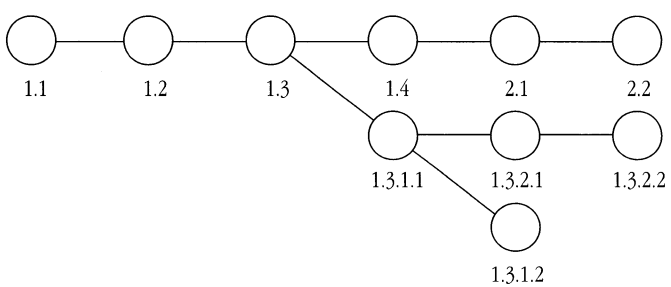

Figure 8-3 A complicated branch structure

If you retrieve version 1.3.2.2, you know that (chronologically) it is the second delta on the second branch from delta 1.3. You are not able to deduce how many deltas there are between version 1.3.2.2 and version 1.3. Thus, although the branching capability is provided for managing files under certain special conditions, it is much easier to manage your files if you keep the SCCS organization as linear and simple as possible.

SCCS command conventions

This section discusses the conventions and rules that apply to SCCS commands. Except where noted, these conventions apply to all SCCS commands. A list of the temporary files generated by various commands (and referred to in the "SCCS Command Summary") is also provided.

SCCS command arguments

SCCS commands accept two types of arguments: keyletters and file arguments.

Keyletters consist of a minus sign followed by a lowercase character, which might be followed by a value. For example, `-a` is a keyletter. Keyletters control the execution of the command to which they are supplied. All keyletters specified for a given command apply to all file arguments of that command. Keyletters are processed before any file arguments, with the result that the placement of keyletters is arbitrary (that is, keyletters can be interspersed with file arguments). Somewhat different argument conventions apply to the `help`, `what`, `sccsdiff`, and `val` commands.

◆ **Note** Keyletters are command-line options equivalent to A/UX flag options. Do not confuse keyletters with SCCS flags, discussed in "SCCS Flags." ◆

File arguments (names of files and/or directories) specify the files to be processed by the given SCCS command. Naming a directory is equivalent to naming all the SCCS files within the directory. Non-SCCS files in the named directories are silently ignored. In general, file arguments cannot begin with a minus sign, but if the name `-` (a single minus sign) is specified as an argument to a command, the command reads the standard input (until end-of-file) and takes each line as the name of an SCCS file to be processed. This feature is often used in pipelines. File arguments are processed left to right.

Flags

Certain actions of SCCS commands can be controlled by flags, which appear in SCCS files. These flags are discussed in "SCCS Flags" later in this chapter.

Diagnostics

SCCS commands produce diagnostics (on the standard error output) that use this format:

`ERROR [`*filename*`] :` *message text* `(`*code*`)`

The code in parentheses can be used as an argument to the `help` command to obtain a further explanation of the diagnostic message. Detection of a fatal error during the processing of a file causes the SCCS command to stop processing that file and to proceed with the next file, in order, if more than one file is named.

Certain SCCS commands check both the *real* and *effective* user IDs (see `passwd`(1) in *A/UX Command Reference*). If you are using SCCS to manage your personal files, these two IDs are the same; if you are working in a group project, see "SCCS Protection Mechanisms" earlier in this chapter.

Temporary files

Several SCCS commands generate temporary files and file copies during the process of creating, retrieving, and updating SCCS files.

The temporary files are normally named by stripping off the `s.` prefix of the SCCS filename and replacing it with another single alphabetic character.

The *g-file* is named by simply deleting the `s.` prefix. Thus, if the SCCS file is named `s.abc` the *g-file* is named `abc`. The *p-file* is named `p.abc`.

Figure 8-4 demonstrates the relationships of the temporary files.

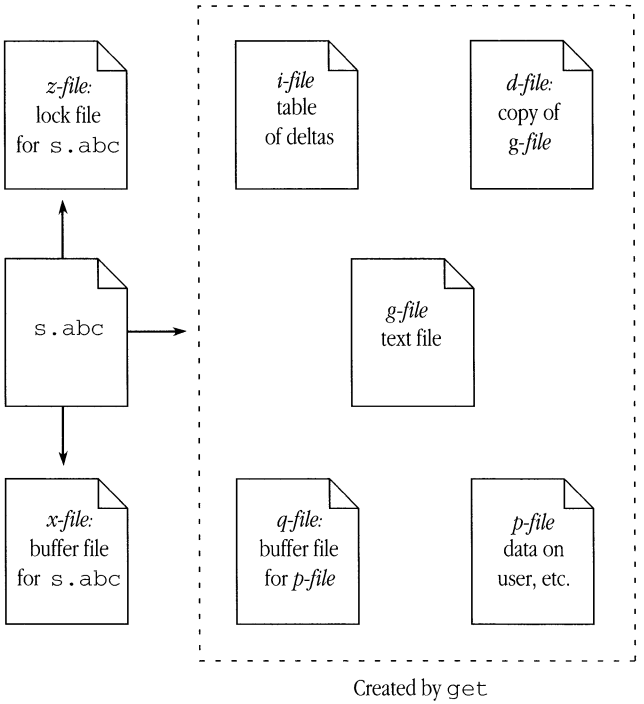

Figure 8-4 Relationships among temporary files

These temporary files are as follows:

d-file

When you invoke a `get` command, SCCS creates its own temporary copy of the *g-file* by performing an internal `get` at the SID specified in the *p-file* entry. This temporary copy is called the *d-file*.

When you record your changes in a new version, the `delta` command compares the *d-file* to to the *g-file* (using the `diff` command). The differences between the *g-file* and the *d-file* are the changes that constitute the delta.

g-file

This is the text file created by a `get` command. It contains a particular version of an SCCS file, and its name is formed by stripping off the `s.` prefix from the SCCS file.

The *g-file* is created in the current directory and is owned by the real user. The mode assigned to the *g-file* depends on how the `get` command is invoked. The version it contains also depends on how the `get` command is invoked. The default version is the most recent trunk delta (that is, excluding branches).

l-file

The `get -l` command creates an *l-file* containing a table showing the deltas used in constructing a particular version of the SCCS file. This file is created in the current directory with mode 444 (read-only) and is owned by the real user.

p-file

When the `get -e` command creates a *g-file* with read-write permission (so you can edit it), it places certain information about the SCCS file (that is, the SID of the retrieved version, the SID to be given to the new delta when it is created, and the login name of the user executing `get`) in another new file called the *p-file*.

When you record your changes in a new version, the `delta` command reads the *p-file* for the SID and the login name of the user creating the new delta.

When the new delta is made, the *p-file* is updated by removing the relevant entry. If there is only a single entry in the *p-file*, then the *p-file* itself is removed.

q-file

Updates to the *p-file* are made to a temporary copy, the *q-file*, whose use is similar to the use of the *x-file*.

x-file

All SCCS commands that modify an SCCS file do so by writing a temporary copy, called the *x-file* (to ensure that the SCCS file is not damaged if processing terminates abnormally). When processing is complete, the old SCCS file is removed and the *x-file* is renamed (with the `s.` prefix) to be the SCCS file.

The *x-file* is created in the directory containing the SCCS file, given the same mode as the SCCS file, and owned by the effective user.

z-file

To prevent simultaneous updates to an SCCS file, commands that modify SCCS files create a "lock file" called the *z-file*. This file exists only for the duration of the execution of the command that creates it. The *z-file* contains the process number of the command that creates it. While the *z-file* exists, it indicates to other commands that the SCCS file is being updated. SCCS commands that modify SCCS files do not process a file if the corresponding *z-file* exists.

The *z-file* is created with read-only mode (mode 444, possibly modified by the user's `umask`) in the directory containing the SCCS file. It is owned by the effective user.

In general, users can ignore most of these temporary files, although they can be useful in the event of system crashes or similar situations.

SCCS ID keywords

When you retrieve an SCCS file to compile it, it is useful to record the date and time of creation, the version retrieved, the module name, and so forth, within the *g-file*. This information appears in a load module when one is eventually created.

SCCS uses ID keywords for recording such information about deltas automatically. ID keywords can appear anywhere in the generated file and are replaced by appropriate values.

The format of an ID keyword is an uppercase letter enclosed by percent signs (`%`). When these appear in the generated SCCS file, they are replaced by the values defined for that keyword. For example,

```
%I%
```

is replaced by the SID of the retrieved version of a file. Similarly,

```
%H%
```

is replaced by the current date (in the form *mm/dd/yy*). When no ID keywords are substituted by `get`, the following message is issued:

```
No id keywords (cm7)
```

This message is normally treated as a warning by `get`, unless the `i` flag is present in the SCCS file (see "SCCS Flags" later in this chapter).

Table 8-1 shows a complete list of the ID keywords.

Table 8-1 SCCS ID keywords

Keyword	Value
`%M%`	Module name: either the value of the `m` flag in the file (see `admin`(1)), or the name of the SCCS file with the leading `s.` removed
`%I%`	SCCS identification (SID) (`%R%.%L%.%B%.%S%`) of the retrieved text
`%R%`	Release
`%L%`	Level
`%B%`	Branch
`%S%`	Sequence
`%D%`	Current date (*yy/mm/dd*)
`%H%`	Current date (*mm/dd/yy*)
`%T%`	Current time (*hh:mm:ss*)
`%E%`	Date newest applied delta was created (*yy/mm/dd*)
`%G%`	Date newest applied delta was created (*mm/dd/yy*)
`%U%`	Time newest applied delta was created (*hh:mm:ss*)
`%Y%`	Module type: the value of the `t` flag in the SCCS file (see `admin`(1))
`%F%`	SCCS file name
`%P%`	Fully qualified SCCS filename
`%Q%`	Value of the `q` flag in the file (see `admin`(1))
`%C%`	Current line number. This keyword is intended for identifying messages sent by the program. It is not intended to be used on every line to provide sequence numbers.
`%Z%`	Four-character string `@(#)` recognizable by `what`
`%W%`	Shorthand notation for constructing `what` strings for A/UX system program files. `%W% = %Z%%M%^I%I%` (where `^I` is the tab character)
`%A%`	Another shorthand notation for constructing `what` strings for non-A/UX system program files. `%A% = %Z%%Y%%M%%I%%Z%`

SCCS command summary

This section describes the features of all the SCCS commands. The SCCS commands are as follows:

`admin`	Creates SCCS files and applies changes to characteristics of SCCS files.
`cdc`	Changes the commentary associated with a delta.
`comb`	Combines two or more consecutive deltas of an SCCS file into a single delta; often reduces the size of the SCCS file.
`delta`	Applies changes (deltas) to the text of SCCS files—that is, creates new versions.
`get`	Retrieves versions of SCCS files.
`unget`	"Undoes" a `get -e` command, if invoked before the new delta is created.
`help`	Prints explanations of diagnostic messages.
`prs`	Prints portions of an SCCS file in user-specified format.
`rmdel`	Removes a delta from an SCCS file; allows the removal of deltas that were created by mistake.
`sact`	Accounts for SCCS files in the process of being changed.
`sccsdiff`	Shows the differences between any two versions of an SCCS file.
`val`	Validates an SCCS file.
`what`	Searches any A/UX system files for all occurrences of a special pattern and prints out what follows it; `what` is useful in finding identifying information inserted by the `get` command.

Create SCCS files: `admin`

`admin` creates new SCCS files or changes characteristics of existing ones. You can create an SCCS file with the command

`admin -i`*filename* `s.`*filename*

where *filename* is a file from which the text of the initial delta of the SCCS file `s.`*filename* is to be taken.

◆ **Note** There is no space between the `-i` keyletter and the *filename* argument. ◆

SCCS files are created in read-only mode (444) and are owned by the effective user (see `passwd`(1) in *A/UX Command Reference*). Only a user with write permission in a directory containing SCCS files can use the `admin` command on a file in that directory.

If you omit the value of the `-i` keyletter, `admin` reads the standard input for the text of the initial delta. Thus, the command

`admin -is.`*filename* `<` *filename*

is also valid. Only one SCCS file can be created at a time using the `-i` keyletter.

If the text of the initial delta does not contain ID keywords, the message

```
No id keywords (cm7)
```

is issued as a warning. See "SCCS ID Keywords" earlier in this chapter for more information.

If you set the `i` flag in the SCCS file (using the `-f` keyletter with the `admin` command; see the next section, "SCCS Flags"), the above message is treated as a fatal error and the SCCS file is not created.

The first delta of an SCCS file is normally 1.1. The `-r` keyletter to the `admin` command is used to specify a different release number for the initial delta. Because it is only meaningful in creating the first delta (with `admin`), its use is permitted only with the `-i` keyletter. The command

`admin -i`*filename* `-r3 s.`*filename*

specifies that the first delta should be named 3.1 rather than 1.1.

SCCS flags

SCCS file flags are used to direct certain actions of SCCS commands.

The flags of an SCCS file are initialized or changed using the `-f` keyletter, and deleted using the `-d` keyletter. When you create an SCCS file, flags are either initialized by the `-f` keyletter on the command line or assigned default values. For example, the following command sets the `i` flag and the `m` (module name) flag:

`admin -i`*filename* `-fi -fm`*modname* `s.`*filename*

The `i` flag specifies that a warning message stating that there are no ID keywords contained in the SCCS file should be treated as a fatal error.

The value *modname* specified for the `m` flag is the value that the `get` command uses to replace the `sccs` ID keyword. (In the absence of the `m` flag, the name of the *g-file* is used as the replacement for the `sccs` ID keyword.)

Note that several `-f` keyletters can be supplied on the `admin` command line and that `-f` keyletters can be supplied whether the command is creating a new SCCS file or processing an existing one.

The `-d` keyletter is used to delete a flag from an SCCS file and can be specified only when processing an existing file. For example, the following command removes the `m` flag from the SCCS file:

`admin -dm s.`*filename*

Several `-d` keyletters can be supplied on a single invocation of `admin` and can be intermixed with `-f` keyletters.

A user list of login names and/or group IDs of users who are allowed to create deltas of that file is checked by several SCCS commands to ensure that the delta is authorized. This list is empty by default, which means that anyone can create deltas. The `-a` keyletter is used to specify users who are given permission or denied permission to create deltas. You can use the `-a` keyletter whether `admin` is creating a new SCCS file or processing an existing one, and it can appear several times on a command line.

For example, the command

`admin -avz -aram -a1234 s.`*filename*

gives permission to create deltas to the login names `vz` and `ram` and the group ID `1234`. The command

`admin -a!vz s.`*filename*

denies permission to create deltas to the login name `vz`. Similarly, the `-e` keyletter is used to remove (erase) login names or group IDs from the list. For example:

`admin -evz s.`*filename*

removes the login name `vz` from the user list of `s.`*filename*.

Comments and MR numbers

When an SCCS file is created, you can insert comments stating your reasons for creating the file. In a controlled environment, it is expected that deltas are created only as a result of some trouble report, change request, trouble ticket, and so forth, all of which are collectively called MRs (for *modification request*).

The creation of an SCCS file might sometimes be the direct result of an MR. MRs can be recorded by number in a delta through the `-m` keyletter, which can be supplied on the `admin` (or `delta`) command line.

The `-y` keyletter can also be used to supply comments on the command line rather than through the standard input. If comments (`-y` keyletter) are omitted, a comment line of the form

`date and time created` *YY/MM/DD hh:mm:ss* `by` *logname*

is automatically generated.

If you want to supply an MR number (using the `-m` keyletter), the `v` flag must also be set (using the `-f` keyletter described below), as in the command

`admin -i`*filename* `-m`*mrlist* `-fv s.`*filename*

The `v` flag causes the `delta` command to prompt for MR numbers as the reason for creating a delta. (See `sccsfile`(4) in *A/UX Programmer's Reference.*) Note that the `-y` and `-m` keyletters are effective only if a new SCCS file is being created.

Descriptive text

The portion of the SCCS file reserved for descriptive text can be initialized or changed using the `-t` keyletter. Descriptive text is intended as a summary of the contents and purpose of the SCCS file.

To insert descriptive text in a file you are creating, the `-t` keyletter is followed by the name of a file from which the descriptive text is to be taken. For example, when a new SCCS file is being created, the following command takes descriptive text from *description-file:*

`admin -i`*filename* `-t`*description-file* `s.`*filename*

When processing an existing SCCS file, the `-t` keyletter specifies that text found in *description-file* should overwrite current descriptive text (if any). If you omit the file name after the `-t` keyletter, as in

`admin -t s.`*filename*

the descriptive text currently in the SCCS file is removed.

Change comments in an SCCS file: `cdc`

`cdc` changes the comments or MR numbers that were supplied when a delta was created. It is invoked as follows:

`cdc -r3.4 s.`*filename*

This specifies that you want to change the comments of delta 3.4 of `s.`*filename*. You can also use `cdc` to delete selected MR numbers by preceding the selected MR numbers by the exclamation character (`!`).

`cdc` prompts for MR numbers and new comments:

```
cdc -r3.4 s.filename
MRs? mrlist! mrlist
comments? deleted wrong MR number and inserted\
correct MR number
```

The new MR numbers in the first `mrlist` are inserted, and the old MR numbers (preceded by the exclamation character) are deleted. The old comments are kept and preceded by a line, indicating that they are changed. The inserted comment line records the login name of the user executing `cdc` and the time of its execution.

Combine deltas to save space: `comb`

The `comb` command generates a shell script (see `sh`(1) in *A/UX Command Reference*) that is written to standard output. When executed, the script attempts to save space by discarding deltas that are no longer useful and combining other specified deltas.

◆ **Note** `comb` should be used only a few times in the life of an SCCS file. Before any actual reconstructions, `comb` should be run with the `-s` keyletter (in addition to any other keyletters desired). ◆

In the absence of any keyletters, `comb` preserves only the most recent deltas and the minimum number of "ancestor" deltas necessary to preserve the SCCS file tree. The effect of this is to eliminate middle deltas on the trunk and on all branches of the tree.

Some of the `comb` keyletters are as follows:

`-p` Specifies the oldest delta that is to be preserved in the reconstruction. All older deltas are discarded.

`-c` Specifies a list of deltas to be preserved (see `get`(1) in *A/UX Command Reference* for the syntax of this list). All other deltas are discarded.

`-s` Causes the generation of a shell script that, when run, produces only a report summarizing the percentage space (if any) to be saved by reconstructing each named SCCS file. You should run `comb` with this keyletter (in addition to any others desired) before any actual reconstructions.

Note that the shell script generated by `comb` is not guaranteed to save space. In fact, it is possible for the reconstructed file to be larger than the original. Note, too, that the shape of the SCCS file tree might be altered by the reconstruction process.

Store a new SCCS file version: `delta`

`delta` creates a new delta by recording the changes made to a *g-file*. The differences between the *g-file* and the *d-file* are the changes that constitute the delta. These changes are normally stored as a delta; they can also be printed on the standard output by using the `-p` keyletter. The format of this output is similar to that produced by `diff`.

Required temporary files

All temporary files used by the `delta` command are described in the previous section, "Temporary Files." There must be a *p-file* and a *d-file* for `delta` to work.

`delta` looks in the *p-file* for the user's login name and a valid SID for the next delta. There should be just one entry for the user (created when the user does a `get -e`) and it should be the same user who is trying to create a delta. Otherwise, `delta` prints an error message and stops. If the user's login name appears in more than one entry in the *p-file*, the same user executed more than one `get -e` on the SCCS file. In this case, the `-r` keyletter must then be used with `delta` to specify the SID that uniquely identifies the *p-file* entry. This entry is the one used to obtain the SID of the `delta` to be created.

The `delta` command also performs the same permission checks performed by `get -e`. If all checks are successful, delta performs a `diff` on the *g-file* and the *d-file* and records the changes as a new delta.

Comments and MR numbers

In practice, the most common use of `delta` is

`delta s.`*filename*

which prompts

```
comments?
```

on the screen. Your response can be up to 512 characters long if you escape all newlines with a backslash (`\`). The response is terminated by a newline character.

In a controlled environment, deltas are usually created only as a result of some trouble report, change request, trouble ticket, and the like. These are collectively called MRs (modification requests) and can be recorded in each delta. If the SCCS file has a `v` flag set, `delta` first prompts with

```
MRs?
```

on the screen. The standard input is then read for MR numbers, separated by blanks and/or tabs. Your response can be up to 512 characters long if you escape all newlines with a backslash (`\`). The response is terminated by a newline character.

The `-y` and/or `-m` keyletters on the `delta` command line can also be used to supply comments and MR numbers, respectively, instead of supplying these through the standard input. The format of the `delta` command is then

`delta -y`*descriptive comment* `-m`*mrlist* `s.`*filename*

The `-m` keyletter is allowed only if the SCCS file has a `v` flag. These keyletters are useful when `delta` is executed from within a shell script (see `sh`(1) in *A/UX Command Reference*).

The `-s` keyletter suppresses all output that is normally directed to the standard output except for the prompts `comments?` and `MRs?`. Use of the `-s` keyletter together with the `-y` keyletter (and possibly the `-m` keyletter) causes `delta` to neither read standard input nor write to standard output.

The comments and/or MR numbers are recorded as part of the entry for the delta being created and apply to all SCCS files processed by the same invocation of `delta`. If `delta` is invoked with more than one file argument and the first file named has a `v` flag, all files named must have the `v` flag. Similarly, if the first file named does not have this flag, then none of the files named can have it. Any file that does not conform to these rules is not processed.

When processing is complete, the SID of the created delta (obtained from the *p-file* entry) and the counts of lines inserted, deleted, and left unchanged by the delta are written to the standard output. Thus, a typical output might be

```
1.4
14 inserted
7 deleted
345 unchanged
```

◆ **Note** The counts of lines reported as inserted, deleted, or unchanged by `delta` might not agree with your perception of the changes applied to the *g-file*. There are usually several ways to describe a set of changes, especially if lines are moved around in the *g-file*, and `delta` is likely to find a description that differs from your perception. However, the total number of lines of the new delta (the number inserted plus the number left unchanged) should agree with the number of lines in the edited *g-file*. ◆

Keywords

If `delta` finds no ID keywords in the edited *g-file*, it prints the message

```
No id keywords (cm7)
```

after it prompts for comments, but before any other output. This indicates that any ID keywords that might have existed in the SCCS file have been replaced by their values or deleted during the editing process. This can be caused by

- creating a delta from a *g-file* that was created by a `get` command without the `-e` keyletter (ID keywords are replaced by `get` in that case)
- accidentally deleting or changing the ID keywords while you are editing the *g-file*
- the file not having any ID keywords to begin with

In any case, it is left up to the user to determine what to do about it. The delta is created whether or not ID keywords are present, unless there is an `i` flag in the SCCS file indicating that this should be treated as a fatal error. In this last case, the delta is not created until the ID keywords are inserted in the *g-file* and the `delta` command is executed again.

See "SCCS ID Keywords" earlier in this chapter for more information.

Removal of temporary files

When processing of an SCCS file is complete, the corresponding *p-file* entry is removed from the *p-file*. All updates to the *p-file* are made to a temporary copy called the *q-file*. If there is only one entry in the *p-file*, then the *p-file* itself is removed.

When processing of the corresponding SCCS file is complete, `delta` also removes the edited *g-file* unless the `-n` keyletter is specified. The command

`delta -n s.`*filename*

keeps the *g-file* upon completion of processing.

Retrieve an SCCS file version: `get`

`get` creates a text file containing a particular version of an SCCS file. The `get` command applies deltas to the initial version of the file to obtain the version you specify or the most recent version (excluding branch versions, which must be retrieved specifically).

The resulting text file is called the *g-file* (see "Temporary Files" earlier in this chapter). The mode of the *g-file* depends on how the `get` command is invoked. For example, the command

`get s.`*filename*

produces

```
1.3
67 lines
No id keywords (cm7)
```

on the standard output. This indicates that version 1.3 (the most recent delta) was retrieved, that there are 67 lines of text in this version, and that no ID keywords were substituted in the file.

The generated *g-file* is assigned mode 444 (read-only), which does not allow you to modify the file, although you can read the file or compile it, and so on. The file is not intended for editing (that is, for making deltas).

If you specify several file arguments (or directory-name arguments) on the `get` command line, similar information is displayed for each file processed, preceded by the SCCS filename. For example, the command

```
get s.abc s.def
```

produces

```
s.abc:
1.3
67 lines
No id keywords (cm7)

s.def:
1.7
85 lines
No id keywords (cm7)
```

See "SCCS ID Keywords" earlier in this chapter.

Retrieving different versions

By default, the `get` command retrieves the most recent delta of the highest-numbered release on the basic trunk of the SCCS file tree (exclusive of branches). To change this default, you can

- Set the `d` flag in the SCCS file. Then, the SID specified as the value of this flag is used as a default.
- Use the `-r` keyletter on the `get` command line to specify which SID you want to retrieve. (If the version you specify does not exist, an error message results.) For example:

  `get -r1.3 s.`*filename*

  In this case, the `d` flag (if any) is ignored. A branch delta can be retrieved similarly:

  `get -r1.5.2.3 s.`*filename*

If you omit the level number

`get -r3 s.`*filename*

the highest-level number (most recent delta) within the given release will be retrieved. If the given release does not exist, `get` retrieves the most recent trunk delta (not in a branch) with the highest-level number within the highest-numbered existing release that is lower than the release you specify.

- Use the `-t` keyletter to retrieve the most recent (top) version in a particular release (when no `-r` keyletter is supplied or when its value is simply a release number). *Most recent* is independent of location in the SCCS tree (see "Delta Numbering" earlier in this chapter). For example, if the most recent delta in release 3 is 3.5,

  `get -r3 -t s.`*filename*

  might produce

  ```
  3.5
  59 lines
  ```

  However, if branch delta 3.2.1.5 were the latest delta (created after delta 3.5), the same command might produce

  ```
  3.2.1.5
  46 lines
  ```

Retrieving a file to create a new delta

When you specify the `-e` keyletter to `get`, the retrieved file has read-write permission and can be edited to make a new delta. For example, the command

`get -e s.`*filename*

produces

```
1.3
new delta 1.4
67 lines
```

on the standard output. The use of `get -e` is restricted (because a new delta can be created), causing a check of the SCCS protection mechanisms (user list and protection flags; see "SCCS Protection Mechanisms" earlier in this chapter). SCCS also checks for permission to make concurrent edits (specified by the `j` flag in the SCCS file; see "Concurrent Edits of Same SID").

If the permission checks succeed, `get -e` creates a *g-file* with mode 644 (readable by everyone, writable only by the owner) in the current directory. This mode can be modified by the user's `umask`.

If a writable *g-file* already exists, `get -e` terminates with an error. This is to prevent inadvertent destruction of a *g-file* that already exists and is being edited for the purpose of making a delta.

ID keywords appearing in the *g-file* are not substituted by `get -e` because the generated *g-file* is subsequently used to create another delta. Replacement of ID keywords causes them to be permanently changed within the SCCS file.

The following keyletters can be used with `get -e`:

`-r` Used to specify a particular version to be retrieved for editing. If the number specified to `-r` does not exist, it is assigned to the new delta.

`-t` Specifies the most recent version in a given release be retrieved for editing.

`-i` Used to specify a list of deltas to be included by `get`. Including a delta means forcing the changes that constitute the particular delta to be included in the retrieved version. This is useful if you want to apply the same changes to more than one version of the SCCS file. When a delta is included, `get` checks for possible interference between those deltas and deltas that are normally used in retrieving the particular version of the SCCS file. Two deltas can interfere, for example, when each one changes the same line of the retrieved *g-file*. Any interference is indicated by a warning that shows the range of lines within the retrieved *g-file* in which the problem might exist. The user is expected to examine the *g-file* to determine whether a problem actually exists and to do whatever is necessary (for example, edit the file). The `-i` keyletter should be used with extreme care.

`-x` Used to specify a list of deltas to be excluded by `get`. Excluding a delta means forcing it not to be applied. This can be used to undo (in the version of the SCCS file to be created) the effects of a previous delta. Whenever deltas are excluded, `get` checks for possible interference between those deltas and deltas that are normally used in retrieving the particular version of the SCCS file. (See the explanation under `-i`.) The `-x` keyletter should be used with extreme care.

`-k` Facilitates regeneration of a *g-file* that might have been accidentally removed or ruined after a `get -e` command, or the simple generation of a *g-file* in which the replacement of ID keywords has been suppressed. A *g-file* generated by the `-k` keyletter is identical to one produced by `get -e`, except that no processing related to the *p-file* takes place (see "Temporary Files" earlier in this chapter).

Concurrent edits of different versions

There is a possibility (in a group project) that several `get -e` commands might be executed at the same time on the same file. However, unless concurrent edits are explicitly allowed (see the next section "Concurrent Edits of Same SID"), no two `get -e` executions can retrieve the same version of an SCCS file. This protection uses information from the *p-file* (see "Temporary Files").

The first execution of `get -e` causes the creation of the *p-file* for the corresponding SCCS file. Subsequent executions only update the *p-file* with a line containing the above information. Before updating, however, `get` checks to ensure that no entry (already in the *p-file*) specifies that the SID (of the version to be retrieved) is already retrieved, unless multiple concurrent edits are allowed. (See the next section, "Concurrent Edits of Same SID.")

If both checks succeed, the user is informed that other deltas are in progress and processing continues. If either check fails, an error message results. It is important to note that the various executions of `get` should be carried out from different directories. Otherwise, only the first execution succeeds because subsequent executions attempt to overwrite a writable *g-file*, which is an SCCS error condition. In practice, such multiple executions are performed by different users so that this problem does not arise (each user normally has a different working directory). (See the section "SCCS Protection Mechanisms" earlier in this chapter for a discussion about how different users are permitted to use SCCS commands on the same files.)

Table 8-2 shows a sample SCCS file retrieved by `get -e` and the SID of the version that is subsequently created by `delta`, as a function of the SID specified to `get`.

In Table 8-2, R, L, B, and S are release, level, branch, and sequence components of the SID. The letter "m" means "maximum." Thus, for example, R.mL means "the maximum level number within release R"; R.L.(mB+1).1 means "the first sequence number on the (maximum branch number plus 1) of level L within release R."

Also note that if the SID specified is of the form R.L, R.L.B, or R.L.B.S, each of the specified components must exist.

The `-b` keyletter is effective only if the `b` flag is present in the file (see `admin`(1)). In this state, an entry of `-i` means "irrelevant."

Table 8-2 Determination of a new SID

SID specified	Keyletter used	Other conditions	SID retrieved	SID of delta to be created
none*	no	R default to mR	mR.mL	mR.(mL+1)
none*	yes	R default to mR	mR.mL	mR.mL.(mB+1)
R	no	R > mR	mR.mL	R.1†
R	no	R == mR	mR.mL	mR.(mL+1)
R	yes	R > mR	mR.mL	mR.mL.(mB+1).1
R	yes	R == mR	mR.mL	mR.mL.(mB+1).1
R	-	R< mR and does not exist	hR.mL ‡	hR.mL.(mB+1).1
R	-	Trunk successor in release >R and R exists	R.mL	R.mL.(mB+1).1
R.L.	no	No trunk successor	R.L	R.(L+1)
R.L.	yes	No trunk successor	R.L	R.L.(mB+1).1
R.L	-	Trunk in release >= R	R.L	R.L.(mB+1).1
R.L.B	no	No branch successor	R.L.B.mS	R.L.B.(mS+1)
R.L.B	yes	No branch successor	R.L.B.mS	R.L.(mB+1).1
R.L.B.S	no	No branch successor	R.L.B.S	R.L.B.(S+1)
R.L B.S	no	No branch successor	R.L.B.S	R.L.(mB+1).1
R.L B.S	-	Branch successor	R.L.B.S	R.L.(mB+1).1

* Applies if the `d` (default SID) flag is not present in the file. If the `d` flag is present in the file, the SID obtained from the `d` flag is interrupted as if it had been specified on the command line. Thus, one of the other cases in this table applies.

† Used to force the creation of the first delta in the new release.

‡ hR is the highest existing release that is lower than the specified, nonexisting, release R.

Concurrent edits of the same SID

Unless the `j` flag is set in the SCCS file (see "SCCS Flags" earlier in this chapter), `get -e` commands are not permitted to occur concurrently on the same SID. That is, `delta` must be executed before another `get -e` is executed on the same SID. If the `j` flag is set in the SCCS file, two or more successive executions of `get -e` on the same SID are allowed. The command

`admin -fj s.`*filename*

sets the `j` flag. Then, the command

`get -e s.`*filename*

might produce

```
1.1
new delta 1.2
5 lines
```

which might be immediately followed by the commands

`mv` *filename new-filename*
`get -e s.`*filename*

The second edit request without an intervening execution of `delta` causes a warning to be generated:

```
1.1
WARNING: being edited: '1.1 1.2 username date-stamp' (ge18)
new delta 1.1.1.1
5 lines
```

In this case, a `delta` command corresponding to the first `get` produces delta 1.2, assuming 1.1 is the latest (most recent) delta, and the `delta` command corresponding to the second `get` produces delta 1.1.1.1.

Keyletters that affect output

The following keyletters affect output:

`-p` The retrieved text is written on standard output rather than on a *g-file*. In this case, all output normally directed to the standard output (such as the SID of the version retrieved and the number of lines retrieved) is directed instead to the standard error output. The `-p` keyletter is used, for example, to create *g-files* with arbitrary names:

`get -p s.`*filename* `>` *filename*

`-s` Suppresses all output that is normally directed to the standard output (the SID of the retrieved version, the number of lines retrieved, and so forth, are not written). This does not affect messages to the standard error output. This keyletter is used to prevent nondiagnostic messages from appearing on the user's terminal, and is often used in conjunction with the `-p` keyletter to pipe the output of `get`. For example:

`get -p -s s.`*filename* `| nroff`

`-g` Suppresses the actual retrieval of the text of a version of the SCCS file. This can be used in a number of ways; for example, to verify the existence of a particular SID in an SCCS file:

`get -g -r4.3 s.`*filename*

This prints the given SID if it exists in the SCCS file or generates an error message if it does not exist. The `-g` keyletter is also used to regenerate a *p-file* that was accidentally destroyed. For example:

`get -e -g s.`*filename*

`-l` Creates an *l-file* named by replacing the `s.` of the SCCS file name with `l`. See "Temporary Files" earlier in this chapter. For example, the command

`get -r2.3 -l s.`*filename*

generates an *l-file* that shows the deltas applied to retrieve version 2.3 of the SCCS file. Specifying a value of `p` with the `-l` keyletter

`get -lp -r2.3 s.`*filename*

causes the generated output to be written to the standard output rather than to the *l-file*. You can use the `-g` keyletter with the `-l` keyletter to suppress the actual retrieval of the text.

`-m` Identifies the changes applied to an SCCS file, line by line. When you specify this keyletter to the `get` command, each line of the generated *g-file* is preceded by the SID of the delta that caused that line to be inserted. The SID is separated from the text of the line by a tab character.

`-n` Causes each line of the generated *g-file* to be preceded by the value of the `%M%` ID keyword (the module name) and a tab character. The `-n` keyletter is most often used in a pipeline with the `grep` command. For example:

`get -p -n -s` *directory* `| grep` *pattern*

searches the latest version of each SCCS file in a directory for all lines that match a given pattern. If both the `-m` and `-n` keyletters are specified, each line of the generated *g-file* is preceded by the value of the `sccs` ID keyword and a tab (caused by the `-n` keyletter) and shown in the format produced by the `-m` keyletter.

Because the contents of the *g-file* are modified when you use the `-m` and/or `-n` keyletters, this *g-file* cannot be used for creating a delta, and neither `-m` nor `-n` can be used with the `-e` keyletter.

Restore a version unchanged: `unget`

If invoked before a `delta`, `unget` undoes a `get -e` command. The following keyletters can be used with `unget`:

`-r`*SID* Uniquely identifies the delta that is no longer intended (the SID for the new delta is included in the *p-file*). This is necessary only if two or more `get -e` commands of the same SCCS file are in progress.

`-s` Suppresses the display of the intended SID of the delta on standard output.

`-n` Retains the *g-file* in the current directory instead of removing it.

For example, the command

`get -e s.`*filename*

followed by

`unget s.`*filename*

causes the last version to be unchanged.

On-line explanations: `help`

The `help` command prints explanations of SCCS commands and the messages printed by some of these commands. If you use `help` without an argument, it prompts for one. Valid arguments are names of SCCS commands or the code numbers that appear in parentheses after SCCS messages. Keyletter arguments or file arguments are not valid arguments to `help`.

Explanatory information related to a command is a synopsis of the command. For example, the command

```
help ge5 rmdel
```

produces

```
ge5:
'nonexistent sid'
The specified sid does not exist in the
given file.
Check for typos.

rmdel:
rmdel -rSID name ...
```

This is printed on standard output by default. If no information is found, `help` prints an error message. Note that `help` processes each argument independently, and an error resulting from one argument will not terminate the processing of the other arguments on the command line.

Print parts of an SCCS file: `prs`

The `prs` command is used to print on the standard output all or parts of an SCCS file in a format you specify. The format is called the output *data specification.* It is a string consisting of SCCS file data keywords (not to be confused with `get` ID keywords), supplied using the `-d` keyletter on the `prs` command line. These keywords can (optionally) be interspersed with text.

Data keywords specify which parts of an SCCS file are to be retrieved and produced. All parts of an SCCS file (see `sccsfile`(4)) have an associated data keyword. Data

keywords are an uppercase character, two uppercase characters, or an uppercase and a lowercase character, enclosed by colons. For example,

```
:I:
```

is the keyword replaced by the SID of a specified delta. Similarly,

```
:F:
```

is the keyword replaced by the SCCS filename currently being processed, and

```
:C:
```

is replaced by the comment line associated with a specified delta. For a complete list of the data keywords, see `prs`(1) in *A/UX Command Reference.*

There is no limit to the number of times a data keyword can appear in a data specification. For example, the command

`prs -d":I: this is the top delta for :F: :I:" s.`*filename*

might produce on the standard output (for example)

`2.1 this is the top delta for s.`*filename* `2.1`

Information can be obtained from a single delta by specifying the SID of that delta using the `-r` keyletter. For example:

`prs -d":F:: :I: comment line is: :C:" -r1.4 s.`*filename*

might produce the following output:

`s.`*filename*`: 1.4 comment line is: THIS IS A COMMENT`

If the `-r` keyletter is not specified, the value of the SID defaults to the most recently created delta.

Information can be obtained from a range of deltas by specifying the `-e` or `-l` keyletters. The `-e` keyletter substitutes data keywords for the SID designated by the `-r` keyletter and all earlier deltas.

`prs -d :I: -r1.4 -e s.`*filename*

might produce

```
1.4
1.3
1.2.1.1
1.2
1.1
```

The `-l` keyletter substitutes data keywords for the SID designated by the `-r` keyletter and all later deltas.

`prs -d :I: -r1.4 -l s.`*filename*

might produce

```
3.3
3.2
3.1
2.2.1.1
2.2
2.1
1.4
```

Substitution of data keywords for all deltas of the SCCS file can be obtained by specifying both the `-e` and `-l` keyletters.

Remove a specific delta: `rmdel -r`

`rmdel` removes a delta from an SCCS file. Normally, you should use it only if incorrect global changes were incorporated in a delta.

The `-r` keyletter is required to specify the complete SID of the delta to be removed.

The delta to be removed must be the most recent delta on its branch or on the trunk of the SCCS file tree. In Figure 8-5, only deltas 1.3.1.2, 1.3.2.2, and 2.2 can be removed; once they are removed, then deltas 1.3.2.1 and 2.1 can be removed.

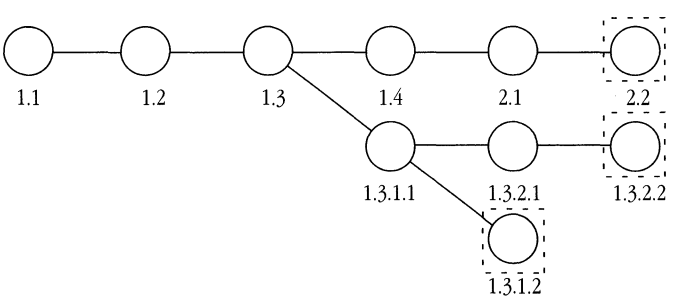

Figure 8-5 Removing a delta

The command

`rmdel -r2.2 s.`*filename*

specifies that delta 2.2 of the SCCS file should be removed. Before removing it, `rmdel` checks that the release number (R) of the given SID satisfies the relation

floor `<=` *R* `<=` *ceiling*

and that the SID specified is not a version that is being changed (for which a `get -e` has been executed and whose associated `delta` has not yet been made).

The A/UX and SCCS protection mechanisms are also checked. If the checks are not successful, processing is terminated and the delta is not removed.

If the checks are successful, the delta is removed and its type indicator in the delta table of the SCCS file is changed from `D` (delta) to `R` (removed).

Account for open SCCS files: `sact`

The `sact` command reports any impending deltas to an SCCS file. An impending delta is a change that has not yet been incorporated into the SCCS file with the `delta` command. This would occur if a `get -e` has been executed but an associated `delta` has not yet been made.

`sact` reports five fields for each named file:

field 1	The SID of the existing SCCS file being changed
field 2	The SID of the new delta to be created
field 3	The login name of the user who executed the `get -e` command
field 4	The date the `get -e` command was executed
field 5	The time the `get -e` command was executed

The command

`sact s.`*filename*

produces a display such as

```
1.2 1.3 john 85/06/20 16:15:15
```

Compare two SCCS files: `sccsdiff`

`sccsdiff` compares two specified versions of one or more SCCS files and prints the differences on standard output. The versions to be compared are specified using the `-r` keyletter in the same format used for the `get` command. For example,

`sccsdiff -r3.4 -r5.6 s.`*filename*

The two versions must be specified as the first two arguments to this command in the order in which they were created (the older version is specified first). Any following keyletters are interpreted as arguments to the `pr` command (which prints the differences on standard output in `diff` format) and must appear before any filenames.

The SCCS files to be processed are named last. Directory names and a name of a single minus sign (`-`) are not acceptable to `sccsdiff`.

Check SCCS file characteristics: `val`

`val` is used to determine whether a file is an SCCS file meeting the characteristics specified by an optional list of keyletter arguments. Any characteristics not met are considered errors.

The `val` command checks for the existence of a particular delta when the SID for that delta is explicitly specified through the `-r` keyletter. The string following the `-y` or `-m` keyletter is used to check the value set by the `t` or `m` flag, respectively (see `admin`(1) in *A/UX Command Reference* for a description of the flags).

The `val` command treats the special argument `-` differently from other SCCS commands. This argument allows `val` to read the argument list from the standard input as opposed to obtaining it from the command line. The standard input is read until an end-of-file.

This capability allows for one invocation of `val` with different values for the keyletter and file arguments. For example:

```
val -
-yc -mabc s.filename
-mxyz -ypl1 s.xyz
(EOF)
```

first checks whether the `s.`*filename* file has a value `c` for its type flag and value *filename* for the module name flag. Once processing of the first file is completed, `val`

then processes the remaining files, in this case `s.xyz`, to determine whether they meet the characteristics specified by the keyletter arguments associated with them.

The `val` command returns an 8-bit code; each bit set indicates the occurrence of a specific error (see `val`(1) for a description of possible errors and the codes). The appropriate diagnostic is also printed unless suppressed by the `-s` keyletter. A return code of zero indicates all named files met the characteristics specified.

Find identifying information: `what`

`what` is used to find identifying information within any A/UX system file whose name is given as an argument to `what`. Directory names and a name of `-` (a single minus sign) are not treated specially as they are by other SCCS commands, and no keyletters are accepted by the command.

The `what` command searches the given files for all occurrences of the string `@(#)` (which is the replacement for the `@(#)` ID keyword) and prints (on the standard output) the balance following that string until the first double quote (`"`), greater than (>), backslash (\), newline, or (nonprinting) null character. For example, if the SCCS file `s.prog.c` (a C language program) contains the following line,

```
char id[] = "@(#)%Z%%M%:%I%";
```

the command

```
get -r3.4 s.prog.c
```

is executed, and the resulting *g-file* is compiled to produce `prog.o` and `a.out`. Then the command

```
what prog.c prog.o a.out
```

produces

```
prog.c:
  prog.c:3.4
prog.o:
  prog.c:3.4
a.out:
  prog.c:3.4
```

The string searched for by `what` does not need to be inserted in the SCCS file through an ID keyword of `get`; it can be inserted in any convenient way.

9 `awk` Programming Language

`awk` operation / 9-3

Comments / 9-5

Command-line options / 9-6

Invocation modes / 9-7

Interactions with the shell / 9-9

Text input processing / 9-11

Patterns / 9-14

Actions / 9-20

Data structures / 9-35

Expressions / 9-41

Lexical conventions / 9-50

Primary expressions / 9-55

Terms / 9-58

Expressions / 9-60

`awk` is a special-purpose language for processing text in terms of input records and fields. The `awk` language can be used to

- generate reports
- match patterns
- tabulate, summarize, and format information
- validate data
- filter data for transmission

Another guide to the `awk` programing language is *The AWK Programming Language* by A.V. Aho, B.W. Kernigan, and P.J. Weinberger (Addison-Wesley, 1988). In addition to its other merits, this book offers fully functional programs for you to use and inspect.

`awk` operation

An `awk` program is a sequence of instructions of the form

pattern { *action* }
pattern { *action* }
...

These pattern-action instructions specify text scanning and text manipulation functions. Sometimes these instructions merely establish settings that affect text processing that is undertaken by `awk` as part of its standard operation.

The standard operation of `awk` is to scan each input file once and look for matches between each input record and any of a set of patterns you supply. An action associated with a pattern is taken while processing each input record that contains text that matches the pattern. So that text patterns can be sought in specific positions in an input record, `awk` automatically splits the input record into fields when it encounters field-separator characters.

After `awk` splits each input record into fields, each field is assigned to a field variable, such as `$1`, `$2`, `$3`, and so forth. These variables can be used to reference input fields either in the pattern or action portions of an instruction. Although `$0` looks like a field reference, it refers to the entire input record with field delimiters unstripped.

When an input record satisfies the pattern criteria, the text of that input record can be accessed through references to the variables `$1`, `$2`, `$3`, and so on (as well as `$0`, the entire input record). For example, to print only those input records containing the string `Mac`, you can use

```
/Mac/ { print $0 }
```

A pattern in front of an action acts as a selector that determines whether that action is performed. After `awk` compares all the patterns to the current input record, it looks at the next input record and repeats the process starting with the first pattern in the `awk` program.

When you want to specify two or more actions for the same pattern, a semicolon or newline character must separate each action. For example, if you want the number 5 printed, you can create the following program:

```
{ x = 5 ; print x }
```

When several actions are performed for a given pattern, they can also be denoted as one **action block.** So

pattern `{` *action* `[;` *action* `]...` `}`

and

pattern `{` *block* `}`

denote the same thing.

Newline characters can also be used to separate actions in an action block:

```
pattern {
  action
  .
  .
  .
  }
```

Occasionally, action blocks are nested inside other action blocks. In such cases, each nested block is delimited with a pair of opening and closing braces:

```
pattern {
  main-block
  {
        sub-block
  }
  }
```

Typically, sub-blocks occur within execution loops or branches, also known as control-flow structures.

In an `awk` program, either the pattern or the action can be omitted, but not both. If there is no action for a pattern, the matching record is simply printed. If there is no pattern for an action, then the action is performed for every input record. (An empty `awk` program does nothing.) Because both patterns and actions are optional, you must enclose actions in braces to distinguish them from patterns. For example, the `awk` program

```
/x/ { print }
```

prints every input record that has the letter `x` in it, as does

```
/x/
```

Some of the possible replacements for *action* are the following statements. Each is discussed later in this chapter. As can be seen, `awk` provides a complete set of flow control constructs.

```
if ( condition ) { block1 } [ else { block2 } ]
while ( condition ) { block }
for ( expression ; condition ; expression ) { block }
break
continue
next
exit
```

Some of the possible replacements for pattern are regular expressions such as `/x/`, and the special patterns, `BEGIN` and `END`. The regular expression syntax is the same as that used by the line editor `ed` as well as other A/UX tools, such as `grep`. The section "Patterns," later in this chapter, describes the pattern possibilities in greater detail.

The `BEGIN` pattern introduces an initialization section that is run before `awk` reads any input records. Likewise, the `END` pattern introduces a finalization section that is run after `awk` exhausts all the input streams. If `awk` encounters `exit` as one of the actions for a pattern, the `END` section is run prematurely and no further input data is read.

You can create variables and assign them values in either the initialization or finalization sections of an `awk` program as well as in the main body. You can assign values from the `awk` command line using parameters (see the upcoming section "Command-line Options"). When you assign the same variable a value in the `BEGIN` section and through parameters on the command line, the command-line assignments are the ones that remain in effect within the main body. The reason for this is that `awk` executes the initialization section before establishing any of the command-line parameters.

Comments

Although comments are neither patterns nor actions, you can include them inside `awk` programs. Comments begin with the character `#` and end with the end-of-line character, as in

```
# this is a comment line
/xyz/    { print "xyz" # this is a comment inside the action}
```

Command-line options

You can use the following arguments for `awk` when you want `awk` program lines to appear in the command line itself:

`awk [-F`*field-separator* `] '`*pattern-action...*`'` [*parameter*]... *input-file...*

You may use the following arguments for `awk` when you maintain the `awk` program lines in a separate file:

`awk [-F`*field-separator* `] -f` *prog-file* [*parameter*]... *input-file...*

If you do not specify input files, `awk` reads from the standard input. Alternatively, `awk` reads from the standard input where you specify hyphen (-) as a filename. The command

```
awk 'program' file1 - file2
```

first reads from `file1`, then from the standard input, and finally from `file2`.

Variables that are initialized on the command line are parameters. Passing shell-maintained values to `awk` is a form of parameter passing. The format of these assignments is similar to variable assignments, except that unescaped spaces cannot be used on either side of the equal sign. (Spaces are treated as argument delimiters, while the entire `awk` parameter must be able to hang together as one command-line argument.)

`awk -f awkfile datafile` *variable1=x variable2=yy*

Often, parameters are used inside a shell script that contains a reference to `awk`. When invoking a shell script and supplying arguments along with its filename, you can pass the argument strings to `awk`. You can reference these arguments from within script command lines as `$1`, `$2`, and so forth, as described in the documentation for each of the shells. In such a case, the `awk` portion of the script might be

`awk -f awkfile` *variable1*`=$1` *variable2*`=$2 datafile`

Normally, you cannot assign values to variables in the `BEGIN` section in this way because the initialization section is evaluated before any command-line *parameters* are supplied. By using the `-v` option in front of an assignment parameter, however, you can make the assigned value available in the `BEGIN` section.

The following example illustrates how `x` can be initialized on the command line as an `awk` parameter.

```
awk '{ print x }' x=5 chap1
```

prints 5 on the standard output once for each input record obtained from the file `chap1`.

To change the field separator, you can use a parameter that makes an assignment directly to the field separator variable `FS`. For example,

```
awk -f awk_program FS=: chap1
```

changes the field separator to a colon. This affects the field-parsing operations of `awk`. (See the section, "`awk` Operation," earlier in this chapter.)

Another way to establish the character to be used as the field separator for field-parsing purposes is to use the `-F` flag option followed by an explicit field-separator character. For example:

```
awk -F: -f awk_program chap1
```

also changes the field separator to the colon character.

The `-F` option can also be followed by a regular expression that specifies one or more characters to be used as field separators. For example:

```
awk -F[:,;] -f awk.program chap1
```

sets the field separator to be any of the three characters comma, semicolon, or colon.

Note that if you specifically set the field separator to a tab (that is, with the `-F` option or by making a direct assignment to `FS`) then `awk` does not recognize blanks as separating fields. If, however, you specifically set the field separator to a blank, tabs are still recognized as separating fields. Certain characters must be escaped to protect them from interpretation by the shell (for example, blank, tab, asterisk, and so forth).

Invocation modes

There are three other ways to invoke `awk` from the command line (brackets appear around optional items).

1. If the program is short, about one or two lines, it is often easiest to specify it directly on the command line:

   `awk` [*flag-options*] '*program*' [*options*] . . .

   where *program* is your `awk` program.

Note that there are single quotes around the contents of the program to prevent the shell from trying to interpret and alter the program. For example, you might enter

```
awk '/findme/' chap1
```

to run the `awk` program consisting of `/findme/` on the input file `chap1`, and obtain a report of all lines containing the string `findme`.

2. Often, it is more convenient to put the program into a separate file, say `awkprog`, and then to tell `awk` to find it from there. To do so, use the `-f` flag option with the `awk` command, as follows:

   `awk -f awkprog [`*other-flag-options*`] ... [`*options*`] ...`

   For example, suppose that you put the following text into a file called `awkprog`:

   ```
   BEGIN {
      print "hello, world"
       exit
   }
   ```

   Then you can give the command

   ```
   awk -f awkprog
   ```

   to the shell, producing

   ```
   hello, world
   ```

   Recall that the word `BEGIN` is a special pattern indicating that the action following in braces is run before any data is read. `print` and `exit` are both discussed in later sections, but their effects here are obvious.

3. Finally, the `awk` program can be put into a file together with the `awk` invocation for use as a shell script of the format

   script-name `[`*script-options*`] ...`

   This becomes handy when the textual input to `awk` needs to be transformed in some static way by other A/UX utilities such as `sort` or `m4`. Sometimes the required preprocessing can also be performed by a first-pass `awk` program, in which case you can stack two calls to `awk` in the same script.

This manner of invocation also allows you to pass command-line arguments to `awk`. After inserting

```
awk '
BEGIN {
   print "hello, '$1' "
   exit
} '
```

in a file called `greet` and establishing execute permission, you can invoke the new shell script from the command line with

```
greet Jim
```

and the output is

```
hello, Jim
```

The next section provides more detailed information about possible interactions with the shell when you present `awk` programs inside of command lines.

Interactions with the shell

Since this is a rather involved topic that has as much to do with shell behavior as with `awk`, you might wish to skip this section. As described in the preceding section, `awk` can be invoked in several ways. You can avoid any possibility of shell interaction with `awk` instructions by using the `-f` flag option as an alternative to presenting an `awk` program inside a command line. This does not prevent you from placing references to the `awk` command in a shell script. However, it does require that a file other than the shell script itself be used to hold the `awk` program instructions.

Sometimes, an `awk` program is contained in a shell script (the `-f` flag option is not used) or is entered interactively within a command line. In either case, the `awk` instructions themselves become subject to processing by the shell, as in

```
awk '/^\.H/ {print "level " $2 " head" }' chap1
```

`awk` interprets many of the same reserved characters as the shell (such as `$` and the double quotation marks in the preceding example). The `awk` program instructions are usually enclosed inside single quotation marks to help ensure that these characters, such as `$2` in the example, are not interpreted by the shell instead of by `awk`. In this way, the shell can be made to pass the `awk` program instructions intact.

Sometimes you might want your `awk` program to interact with the shell, so you deliberately place `awk` instructions inside a shell script, and perhaps alter the manner of escape (from single quotation marks to double quotation marks, for example) to allow the shell to interpret the `awk` instructions. This can become complicated because of the similarity of variable names built into the shell and into `awk`. It might take some work to get parameters passed from the shell into `awk` program lines.

Suppose you want to write an `awk` program to print lines containing the text specified as an argument to the script. That is, you want a program called `search` so that

```
search Macintosh chap1
```

runs the `awk` program

```
awk '/Macintosh/ { print }' chap1
```

How does the value Macintosh get from the command line into the `awk` program? There are several ways to do this. One is to define `search` as a shell script, as follows:

```
# search: print each record containing the
# string specified in the first argument
# inside the file that is the second argument
awk '/'$1'/ { print }' $2
```

◆ **Note** These are not nested quotes. ◆

When the shell parses this script, it does not interpret anything contained within the first pair of single quotes (`'/'`), but passes it as input to `awk`. Because it is outside of the protective quotes, `$1` is acted upon by the shell, which replaces it with the first argument given on the command line. The shell then passes the remaining portion of the `awk` program without attempting to interpret any of it, since it is enclosed in single quotes.

Note that the string passed in the first argument to the script must not contain a space character (which is a little unusual because unescaped spaces are commonly used to delimit arguments). If you did use an escaped space in the first argument, `awk` would "see" an incomplete program and it would also "see" the last part of the program as if it were a discrete input file argument—causing it to try to open a nonexistent file. To avoid this problem, consider using double quotation marks around the first argument (`$1`). For example:

```
'/' "$1" '/ '
```

In general, the escape character used to set off the `awk` instructions from the rest of the command line becomes more difficult, if not impossible, to use inside the program.

Text input processing

The default behavior of `awk` is for the `awk` instructions in your programs to have a chance to execute once for each input record read. This section describes the default behavior of `awk` in greater detail and tells you how to alter it.

The way `awk` determines that it has read sufficient characters from the input source to obtain one complete record is by scanning the input for one or more record-separator characters. By default, the end-of-record character is a newline. So, an input record normally corresponds to a single line from the input source. However, if the end-of-file character is reached without an immediately preceding end-of-record character, then the end-of-file character is treated as if it were also an end-of-record character.

Accordingly, a record is a sequence of characters from the input ending with a newline character or with an end-of-file character. You can change the character that indicates the end of a record by assigning a new character to the special variable `RS` (the *record separator* variable).

Once `awk` reads a record, it splits the record into *fields*, determined through occurrences of one or more field-separator characters. By default, the end-of-field character is a space or tab character. Accordingly, a field is a sequence of characters derived from the input record that does not contain blanks or tabs. You can change the field-separator character by assigning a new character to the special variable `FS` (the *field separator* variable).

To help explain the text input processing that `awk` normally performs, the precise composition of a sample file is described next, and many examples throughout the remainder of this chapter show what happens to this sample file after processing by an `awk` program.

For explanatory purposes, assume that a file named `countries` has been created and that it contains information about the ten largest countries in the world, including the area in thousands of square miles, the population in millions, and the continent. (Figures are from 1978; `Russia` is placed in `Asia`.) The `countries` file looks like this before it is processed:

```
Russia      8650      262      Asia
Canada      3852       24      N. America
China       3692      866      Asia
USA         3615      219      N. America
Brazil      3286      116      S. America
Australia   2968       14      Australia
India       1269      637      Asia
Argentina   1072       26      S. America
Sudan        968       19      Africa
Algeria      920       18      Africa
```

The wide spaces are tabs in the original input, while a single blank separates `N.` and `S.` from `America`. This sample file is used as the input for many of the `awk` sample programs in this guide because it is typical of the kind of material that `awk` is best at processing (a mixture of words and numbers separated into fields or columns separated by blanks and tabs).

Each of the lines in the sample file has either four or five fields if the default field separators are not altered. So, using the default settings, the first record of the countries file that `awk` parses is

```
Russia 8650 262 Asia
```

Once parsed, this text (less the newline) is assigned to the variable `$0`. If you want to refer to the full text of an input record that `awk` is processing, use the variable `$0`. For example, the following action:

```
{ print $0 }
```

prints the entire record.

Once parsed, the fields within an input record are stored in the variables `$1`, `$2`, `$3`, and so forth. Use `$1` to refer to the first field, `$2` to refer to the second field, `$3` to refer to the third field, and so forth. If you refer to a field number that is higher than the field that was last parsed, you reference an empty string.

Once records and fields are parsed, `awk` also sets certain variables that provide additional information about the current state. These are the built-in variables that are used to maintain record and field counts:

`NF`	the number of fields parsed from the current input record
`NR`	the total number of records fetched so far
`FNR`	the number of input lines fetched with respect to the current input file

The variable `FILENAME` is set to the name of the current source of input. Thus, while `awk` processes the first record of the file countries, `$1` is equal to the string `Russia`, `$2` is equal to the string `8650`, `NF` is equal to 4, `FILENAME` is equal to `countries`, `NR` is equal to 1, and `FNR` is equal to 1.

The following examples show different ways to take advantage of the default text input processing of `awk`.

To print the number of fields, followed by the continent, the name of the country, and the country's population, run the following `awk` program:

```
awk '{ print NF, $4, $5, $1, $3 }' countries
```

to produce

```
4 Asia Russia 262
5 N. America Canada 24
4 Asia China 866
5 N. America USA 219
5 S. America Brazil 116
4 Australia Australia 14
4 Asia India 637
5 S. America Argentina 26
4 Africa Sudan 19
4 Africa Algeria 18
```

Note that values referenced by `$4` and `$5` together comprise the continent name because `N. America` and `S. America` contain an embedded space.

To produce a numbered list of the values in the first field in the file countries, you can enter

```
awk '{ print NR, $1 }' countries
```

which prints the following line numbers and names of countries:

```
1 Russia
2 Canada
3 China
4 USA
5 Brazil
6 Australia
7 India
8 Argentina
9 Sudan
10 Algeria
```

Patterns

A pattern in front of an action acts as a selector that determines whether the action is to be executed. Several types of expressions can be used as patterns:

- the special patterns `BEGIN` and `END`
- isolated regular expressions
- expressions that evaluate to true or false

An isolated regular expression also evaluates to true or false, so the second and third pattern cases are essentially the same. (Only the second case does not qualify as a syntactially complete expression, since it lacks an operator.)

Expressions are used more frequently as patterns than are the `BEGIN` and `END` patterns.

Using expressions for patterns

The pattern portion of an `awk` instruction is equivalent to a conditional expression, since it must produce only a true or false result that gates an action. See the section "Expressions" later in this chapter for a lengthier explanation regarding expressions, including conditional expressions.

By using the operators `~` or `!~`, you can create patterns that return true when the value of a field or variable is sought by a regular expression, or when its value is not sought by a regular expression. This form of operation is either called a *matching operation* or a *pattern-seeking operation.* In general, the format of this type of operation is as follows:

string ~/*pattern*/

Such an expression returns true if the string given in place of *string* contains a substring that is sought by the regular expression *pattern.* For example:

```
$0 ~ /Total/
```

evaluates to true if the current input record contains the string `Total`. With the input file `countries` as before, the program

```
$1 ~ /ia$/ {print $1}
```

prints all countries (field 1 items) whose names end in `ia` (The `$` symbol represents the end of a string. An explanation of the symbol can be found in the next section):

```
Russia
Australia
India
Algeria
```

As it turns out, the following two patterns are equivalent:

```
$0 ~ /Total/
/Total/
```

So, whenever an isolated regular expression is specified, `awk` performs this pattern-seeking operation, yielding a true result whenever `awk` finds at least one string sought by the regular expression somewhere within the current input record.

To be able to construct a regular expression that can seek many different substrings, refer to the next section, "Regular Expression Syntax," for more detailed information.

To construct a pattern that returns true for a contiguous set of input records starting from the first record that matches *pattern1* and ending with the first record that matches *pattern2*, specify two regular expressions separated by a comma:

pattern1, *pattern2*

Any expression that returns true or false upon evaluation is suitable for use as a pattern. So, expressions consisting of comparisons between strings of characters or numbers can be an `awk` pattern. For example, if you want to print only countries with more than 100 million population, use

```
$3 > 100
```

This tiny `awk` program is a pattern without an action, so it prints each line whose third field is greater than 100, as follows:

```
Russia 8650 262 Asia
China 3692 866 Asia
USA 3615 219 N. America
Brazil 3286 116 S. America
India 1269 637 Asia
```

To print the names of the countries that are in Asia, type

```
$4 == "Asia" {print $1}
```

which produces

```
Russia
China
India
```

A comparison expression is one that makes use of a comparison operator, such as

```
<
<=
==
!=
>=
>
```

In such comparisons, if both operands are numeric, a numeric comparison is made. Otherwise, the operands are compared as strings. Thus,

```
$1 >= "S"
```

selects lines that begin with `S`, `T`, `U`, and so forth, which in this case is

```
USA 3615 219 N. America
Sudan 968 19 Africa
```

In the absence of other information, fields are treated as strings, so the program

```
$1 == $4
```

compares the first and fourth fields as strings of characters and prints the single line

```
Australia 2968 14 Australia
```

If the variables for two fields contain numbers, comparisons involving two such field variables are performed numerically.

Regular expression syntax

These additional search capabilities make use of regular expressions. The simplest regular expression is a literal string of characters enclosed in slashes.

```
/Asia/
```

This is a complete `awk` program that prints all lines that contain any occurrence of the string `Asia`. If a line contains `Asia` as part of a larger word like `Asiatic`, the larger word is also printed (but there are no such words in the `countries` file).

`awk` regular expressions are often like those found in the text editor `ed` and the pattern finder `egrep`, in which certain characters have special meanings. For example, you can print all lines that begin with `A` by using

```
/^A/
```

or all lines that begin with `A`, `B`, or `C` by using

```
/^[ABC]/
```

or all lines that end with `ia` by using

```
/ia$/
```

The circumflex (^) means "match the beginning of a line." The dollar sign ($) means "match the end of the line," and enclosing characters in brackets ([and]) means "match any of the characters enclosed." In addition, awk allows parentheses for grouping, the vertical bar (|) for alternatives, the plus sign (+) for "one or more" occurrences, and the question mark (?) for "zero or one" occurrences. For example:

```
/x|y/ { print }
```

prints all records that contain either an x or a y. And

```
/ax+b/ { print }
```

prints all records that contain an a followed by one or more x characters followed by a b. For example: axb, Paxxxxxxxb, QaxxbR.

```
/ax?b/ {print}
```

prints all records that contain an a followed by zero or one x followed by a b. For example: ab, axb, yaxbPPP, CabD.

The two characters . and * have the same meaning as they have in ed or grep: namely, . matches any character and * matches zero or more occurrences of the character preceding it. For example:

```
/a.b/
```

matches any record that contains the letter a followed by any character followed by the letter b. That is, the record must contain an a and a b separated by exactly one character. For example, /a.b/ matches axb, aPb, and xxxxaXbxx, but not ab or axxb.

```
/ab*c/
```

matches a record that contains an a followed by zero or more b characters followed by a c. For example, it matches ac, abc, and pqrabbbbbbbbbbc901.

It is possible to turn off the special meaning of metacharacters such as ^ and * by preceding these characters with a backslash. An example of this is the pattern

```
/\/.*\//
```

which matches any string of characters enclosed in slashes.

BEGIN and END

The awk program executes the action corresponding to the special pattern BEGIN before the input is read. The action corresponding to the special pattern END is executed after all the input is processed. Thus, BEGIN provides a way to gain control before processing for initialization and END helps you wrap up after processing.

You can use BEGIN to put column headings on the output. For example, if you put the following awk program in the file awkprog:

```
BEGIN
{ print "Country",
        "Area",
        "Population",
        "Continent"
         }
       { print }
```

and invoke awk with the command line

```
awk -f awkprog countries
```

The output is

```
Country Area Population Continent
Russia 8650 262 Asia
Canada 3852 24 N. America
China 3692 866 Asia
USA 3615 219 N. America
Brazil 3286 116 S. America
Australia 2986 14 Australia
India 1269 637 Asia
Argentina 1072 26 South Africa
Sudan 968 19 Africa
Algeria 920 18 Africa
```

◆ **Note** Formatting is obviously not very good here; printf does a better job and is usually mandatory if you really care about appearance (see the section "printf" for more information). ◆

Recall also that the `BEGIN` section is a good place to establish settings for special variables, such as `FS` or `RS`, that affect the record-parsing and field-parsing activity about to occur. For example,

```
BEGIN
{ FS = "       "    #tab
 print "Country",
 "Area",
 "Population",
 "Continent"
 }
 { print }
END { print "The number of records is", NR }
```

contains an initialization section that assigns a tab to `FS` for use as the field separator. As a result, all records (in the file `countries`) have exactly four fields.

Actions

The most common type of action is the evaluation of an expression. The operators and functions that help form expressions perform most of your computational work. However, expressions are allowable only in certain action and control-flow contexts. This section explains where expressions can occur within the larger context of actions. "Expressions," a separate section later in this chapter, covers the specification of text-manipulating and number-computing expressions in detail.

The simplest actions are probably those that print a constant string or a number value such as

```
{ print "Hello world" }
```

Other output and input functions that are available are

```
print
printf
getline
```

Actions that assign discrete values to a variable or an array are about as simple as printing statements. Note that there is no need to declare the name of a variable or its data type in advance of its use:

variable = *value*

More detailed information about variables and arrays is provided in the "Data Structures" section later in this chapter.

Two topics that should also be considered along with input and output functions are the redirection of input and the redirection of output. Because these are more advanced topics, they are discussed last.

Components of `awk` programs

Input, output, and assignment actions can be placed nearly anywhere inside the action portion of a pattern-action instruction.

Other action statements must be placed at appropriate locations. For example, certain flow-control actions, such as `break` and `continue`, make sense only inside a looping construct.

Precisely speaking, the action placeholder represents one or more action statements.

Control-flow structures require introductory keywords, such as `while` and `if`, and require a particular ordering of these and other elements:

`if` (*condition*) { *block1* } [`else` { *block2* }]

`while` (*condition*) { *block* }

`for` (*expression* ; *condition* ; *expression*) { *block* }

An *expression* is a sequence of values (or subexpressions that evaluate to values) interspersed with the operations that are performed on them. Refer to "Expressions" later in this chapter for more detailed information concerning expressions. Refer to "Flow of Control" later in this chapter for a more detailed explanation of each of the control-flow structures.

Operator symbols such as

```
+
-
*
%
/
```

can appear throughout long expressions. These can be categorized as binary operators since they act upon values (or subexpression results) to their left and their right:

left-value binary-operator right-value

To help denote nested subexpressions in the following listing, *binary-expression* represents the binary operation structure. It can appear inside of an enclosing binary expression as the left operand, the right operand, or both operands:

binary-expression	*binary-operator*	*right-value*
left-value	*binary-operator*	*binary-expression*
binary-expression	*binary-expression*	*binary-operator*

Parentheses can be used to establish the order of evaluation when the default operator precedence is not desired.

Other basic ways to transform a value involve functions. Each function affects a particular number of input values (individual function syntax formats are about to be supplied), and evaluates a single string or number result based upon them. The overall format for a function is

function-name (*value* [*value-sep* *value*] . . .)

Wherever a value placeholder appears in the function or binary-operation syntax, either an expression (even a function-containing one) or a function can be supplied:

function-name (*expression* [*value-sep* *expression*] . . .)
function (*values*) *binary-operator* *function* (*values*)

A **condition** is any expression that evaluates to either true or false. Comparison operators can be used as binary operators inside of expressions to obtain a true or false result. In such a case, the left or right value for the comparison operator can also be a subexpression or function. In the following example, the left operand for the greater-than operation is an expression, the result of which is compared to a variable named `max`:

```
$1 % 9 > max
```

Certain flow-control constructs, such as `if` and `while` structures, require conditions. Such structures are designed to cause certain actions to be skipped when the result of a condition is false. The *pattern* portion of an `awk` must evaluate to true or false to select the associated action or actions. So a *pattern* is really one form of a condition expression, and can even be technically considered a flow-control construct.

Although it can be done, using a variable assignment as the left or right operand for a comparison makes little sense. The value "returned" by an assignment is a boolean value (always equal to "true"). So, composing a condition based on an assignment really makes the performance of the associated action static rather than conditional.

Flow of control

Besides the flow control established through the associations between patterns and actions, more traditional flow-control constructs are available in the action component.

The control structures for the `awk` language are

- `if-else`
- `while`
- `for`

They are used to establish the flow of evaluation of actions based on the value resulting from a conditional expression. For looping constructs, the condition is evaluated repeatedly until a loop-terminating value is reached.

The conditional expression can include subexpressions, as long as the result finally evaluated is a true or false condition. For example, it can include regular expressions that are specified along with one of the "pattern-seeking" (`~` and `!~`) operators. To test multiple true and false conditions, use the logical operators as well. (See "Combining true-or-false expressions" later in this chapter.) Finally, it can include parentheses for grouping.

A more complete treatment of these structures is given throughout the remainder of this section. In general, the syntax formats for these structures closely follow the corresponding control structures of the C language.

For the looping constructs, the flow-altering functions `break` and `continue` are available. Use `break` to terminate further loop iterations and to skip past any remaining code in the current loop iteration. Use `continue` to skip past any remaining code in the current loop iteration, then continue with the next iteration.

In addition to the control-flow structures, there are the flow-establishing statements `next` and `exit`.

The `next` statement skips past any remaining lines of `awk` instructions for the current input record, finds the next input record, and resumes processing from the beginning of the `awk` program. (Note the difference between `next` and `getline`. `getline` does not skip to the top of the `awk` program.)

An `exit` statement in the `BEGIN` section of an `awk` program stops further program execution so even the `END` section (if there is one) is not executed. An `exit` statement in the main body of the `awk` program stops execution of the main body of the `awk` program. No more input records are reviewed, but the `END` section is executed. An `exit` statement in the `END` section causes execution to terminate at that point.

The remainder of this section provides the syntax description for the three major control structures.

The `if` statement is used as follows:

`if (` *condition* `) {` *block1* `} [else {` *block2* `}]`

The *condition* is evaluated; if it is true, *block1* is executed. Otherwise, *block2* is executed. The `else` part is optional. In the context of an `if` construct, any number of actions enclosed in braces (`{ }`) are either evaluated or skipped as a block, depending on the value resulting from *condition*. One way to determine the country with the maximum population using an `if` construct is

```
{
        if (maxpop < $3)
        {
                maxpop = $3
                country = $1
        }
}

END { print country, maxpop }
```

The `while` loop syntax is:

`while (` *condition* `)` *block*

The *condition* is evaluated; if it is true, the *block* is executed. The *condition* is evaluated again, and if true, the *block* is executed. The cycle repeats as long as the *condition* is true. For example, the following action prints all input fields one per line:

```
{
i = 1
      while (i <= NF)
      {
            print $i
            i++
      }
}
```

Another example is the Euclidean algorithm for finding the greatest common divisor of two values:

```
{ print "the greatest common divisor of"
      print $1 " and " $2 " is "
      while ($1 != $2)
      {
            if ($1 > $2) $1 = $1 - $2
            else $2 = $2 - $1
      }
  print $1
}
```

The `for` loop syntax is similar to that of C.

`for (` *expression1* `;` *condition* `;` *expression2* `)` *block*

This has the same effect as

expression1

`while (` *condition* `)`

```
 {
 statement
 .
 .
 .
 expression2
 }
```

So,

```
{ for (i=1 ; i <= NF; i++) print $i }
```

is another `awk` program that prints all input fields one per line. Note that multiple initializations are not permitted, as in

```
for (i=1,j=2; .\|.\|.; .\|.\|.)
```

The alternative form of the `for` loop is suitable for accessing the elements of an array:

`for (` *var* `in` *array* `)` *block*

performs *block* once for each element in the array after assigning the subscript used to access the element to the variable *var*. The subscripts are accessed in no predictable order. Chaos ensues if the variable *var* is altered or if any new elements for *array* are assigned within the loop.

You can use this form of the `for` loop to print each input record preceded by its record number (`NR`):

```
 { x[NR] = $0 }
END { for(i in x) { print i, x[i] } }
```

A more practical example is the following use of strings to index arrays to add the populations of countries by continents:

```
BEGIN { FS="\t" }
 { population[$4] += $3 }
END {
for (i in population)
print i, population[i]
}
```

In this program, the body of the `for` loop is executed for `i` equal to the string `Asia`, then for `i` equal to the string `N. America`, and so forth, until all the possible values of `i` are exhausted; that is, the program is repeated until all the strings of names of continents are used. Note, however, that the order in which the loops are executed is not specified. If the iteration associated with `N. America` is executed before the iteration associated with the string `Asia`, such a program might produce the following:

```
S. America 142
Africa 37
N. America 243
Asia 1765
Australia 14
```

Report generation

The flow-of-control statements in the last section are especially useful when `awk` is used as a report generator. `awk` is useful for tabulating, summarizing, and formatting information. The last section shows an example of `awk` tabulating populations. Following is another example of this. Suppose you have a file, `prog.usage`, that contains lines of three fields: `name`, `program`, and `usage`. For example:

```
Smith draw 3
Brown eqn 1
Jones nroff 4
Smith nroff 1
Jones spell 5
Brown spell 9
Smith draw 6
```

The first line indicates that `Smith` used the `draw` program three times. If you want to create a program that has the names in alphabetical order and then shows the total usage, use the following program, called `list.a`:

```
{ use[$1 "\t" $2] += $3 }
END {
for (np in use)
print np "\t" use[np] | "sort +0 +2nr"
}
```

This program produces the following output when used on the input file `prog.usage`:

```
Brown eqn 1
Brown spell 9
Jones nroff 4
Jones spell 5
Smith draw 9
Smith nroff 1
```

If you want to format the previous output so that each name is printed only once, pipe the output of the previous `awk` program into the following program, called `format.a`:

```
{
if ($1 != prev)
{
      print $1 ":"
      prev = $1
}
print "\t" $2 "\t" $3
}
```

The variable `prev` prints the unique values of the first field. The command

```
awk -f list.a prog.usage | awk -f format.a
```

gives the output

```
Brown:
eqn   1
spell 9
Jones:
nroff 4
spell 5
Smith:
draw  9
nroff 1
```

It is often useful to combine different `awk` scripts and other shell commands, such as `sort`, as was done in the `list.a` script on the preceding page.

Reading input: `getline`

The `getline` function instructs `awk` to read the next input record, despite the fact that many pattern-actions might not get a chance to execute for the preceding input record. Furthermore, control is left at exactly the same spot in the `awk` program, rather than resuming at the start of the program, as with `next` (see "Flow of Control" earlier in this chapter).

Whether the field-parsing functions previously discussed (see "Text Input Processing") are performed, depends on whether `getline` is specified with a variable name as an argument. If a variable reference is present they are skipped, leaving it up to you to specify a particular field-parsing function (see the description of the split in "Built-in String Functions").

Here are the forms you can use:

`getline`

`getline <`*file*

`getline` *variable*

`getline` *variable* `<`*file*

For the first form, field-referencing variables such as `$0`, `$1`, and so on, are all set, as well as the field and record-counting variables `NR` and `FNR` and `NF`. The second form does not increment the record-counting variables but it does set the field-counting variable (`NF`).

The third and fourth forms shown do not parse the input line into fields, but they do read it into the named variable. These forms of the command also do not set any of the field-referencing variables and do not set the field-counting variable. The third form increments the record-counting variables, but the fourth form does not. In short, the fourth form affects none of the built-in variables.

Two forms of the command involve input redirection: The files named as sources of input on the command line are ignored. Instead, the file to be read from is the one that is supplied as an argument following the `<` redirection symbol within the `getline` statement. In these cases, the getline function returns 0 for the end-of-file character and 1 for a normal record. A handy use for these forms of the `getline` statement is the initialization of array elements in the `BEGIN` section of a program, as in the following example.

```
BEGIN {
        count = 1
        while ( getline array[count] <"table" > 0 )
        { count = count + 1 }
 .
 .
 .
```

A similar example follows. It uses the first field of each record as the subscript for an array element and the second field as the value to be assigned to the subscripted array element.

```
BEGIN {
count = 1
 while ( getline <"table" > 0 )
{ array[$1] = $2 }
 .
 .
 .
```

Note that there is an upper limit to the number of files that can be read this way. However, through use of the `close` function, you can work with an indefinite number of files as long as you don't try to keep them all open at once. The syntax for `close` is

`close(`*file*`)`

For related discussions, see "Directing Output to Other Programs" later in this chapter.

Printing output: `print` and `printf`

The output functions include two forms of print statements, including one that resembles the C function used for printing. Either one can be used, but the C-like `printf` function is capable of formatting its arguments however you want, such as in a dollars-and-cents format.

The print statements normally guide data to the standard output, but both forms of the print command also allow redirection into a file named within `awk` statements:

print-command `[`*expression-list*`]` `[>`*file*`]`

If the redirection symbol `>` is replaced by `>>`, output is appended to the file rather than overwriting it.

Use quotation marks around *file* if *file* is not a string constant. Without quotation marks, the filenames are likely to be treated as variables that, upon reference, are initialized to empty strings.

So, besides redirecting all output on the command line used to invoke `awk`, you have the option within your program to write individual items of data into specific files.

Using the previous example, with the input file countries, you might want to print all the data from countries in `Asia` in a file called `ASIA`, all the data from countries in `Africa` in a file called `AFRICA`, and so forth. To do so, use the following `awk` program:

```
$4 ~ "Asia" { print > "ASIA" }
$4 ~ "Europe" { print > "EUROPE" }
$4 ~ "North" { print > "N_AMERICA" }
$4 ~ "South" { print > "S_AMERICA" }
$4 ~ "Australia" { print > "AUSTRALIA" }
$4 ~ "Africa" { print > "AFRICA" }
```

Note that there is an upper limit to the number of files that are written in this way. However, through use of the `close` function, you can work with an indefinite number of files as long as you don't try to keep them all open at once.

In general, you can direct output into a file after a `print` or a `printf` statement by using a statement of the form

print `>` "*file*"

where *file* is the name of the file receiving the data, and the *print* statement can have any of its allowable arguments.

print

The overall format for the `print` command is

`print` [*expression*] . . . [>*file*]

One of the simplest actions is to print each line of the input to the output, which can be performed by using `print` without a specified pattern. An action might not have a pattern, and in this case `awk` executes the action for all of the lines and prints the entire input record.

```
{ print }
```

To print an empty line, use

```
print ""
```

To print one or more fields in the current input record, replace *expression* with references to field variables. For instance, when using the previously described file countries for data input, the command line

```
awk '{ print $1, $3 }' countries
```

prints the names of the countries and their populations:

```
Russia 262
Canada 24
China 866
USA 219
Brazil 116
Australia 14
India 637
Argentina 26
Sudan 19
Algeria 18
```

There are two special variables that affect the `print` command, `OFS` and `ORS`.

Items (expressions) that are separated by commas within the `print` statement are regarded as fields, so the `print` statement inserts the character that is established as the output field separator between them. By default, the output field separator is a space. The output field separator (`OFS`) is a variable.

The value stored in `ORS` is the output record separator, which `awk` places at the end of any (evaluated) expressions. By default, the output record separator is the newline character. The output field separator (`ORS`) is a variable.

In the following example,

```
{ x="hello"; y="world"; print x, y ; print y x }
```

the default field separator (blank) is used in the first print statement, but not the second, producing

```
hello world
worldhello
```

To place a comma within the output, you can either insert it in the print statement, as in this case:

```
{ x="hello"; y="world" ; print x "," y }
```

or you can change `OFS` in the `BEGIN` section, as in

```
BEGIN { OFS=", "}
 { x="hello"; y="world" ; print x, y }
```

Both of these last two programs yield

```
hello, world
```

printf

For more demanding printing problems, `awk` also provides a C-like `printf` statement. Before printing, `printf` formats strings or numbers in accordance with *format-string*, as the following syntax description shows:

`printf` *format-string*, *expr* `[,` *expr* `]...`

The *format-string* format specifier is exactly like the one used with `printf` in the C library, except that the formatting symbol `*` is not supported. For example,

```
{ printf "%10s %6d %6d\n", $1, $2, $3 }
```

prints `$1` as a string of ten characters (right-justified). The second and third fields (six-digit numbers) make a neatly columned table:

```
Russia       8650        262
Canada       3852        244
China        3692        866
USA          3615        219
Brazil       3286        116
Australia    2968         14
India        1269        637
Argentina    1072         26
Sudan         968         19
Algeria       920         18
```

With `printf`, no output separators or newlines are produced automatically. You must add them, as in this example. As in the C library version of `printf`, the escape characters `\n` (newline) and `\t` (tab) are valid with `printf`.

The `system` command

The `system()` command allows you to run another program, in fact, any UNIX command, from inside an `awk` script. The command has the format

`system(`*expression*`)`

where *expression* is a string. The `system` command executes the command in *expression*. For example:

```
system("cat " $1)
```

runs the `cat` command on the file whose name is in the first field of the input line. The command

```
system("date")
```

runs the `date` command.

Output can be created by using the `system()` command, if the command used as *expression* creates the output.

Directing output to other programs

It is also possible to direct printing into a pipe instead of a file. For example:

```
{ if ($2 == "XX") print | "mail harry" }
```

(where `harry` is a login name), any record with the second field equal to `XX` is sent to the user `harry` as mail. But instead of passing each such record across the pipe to mail individually, `awk` waits until the entire print input is processed before passing its output on to mail. Also,

```
{ print $1 | "sort" }
```

takes the first field of each input record, accumulates them until the input to print is exhausted, and then passes the entire list to sort, which then generates the sorted list. The command in double quotation marks can be any A/UX command.

Only one output pipe is permitted in an `awk` program at one time.

However, through use of the `close` function, you can work with an indefinite number of pipes as long as you don't try to keep them all open at once. If you want to write a file and then read it later, you must close it in between.

In all output statements involving redirection of output, the files or pipes are identified by their names, but they are created and opened only once in the entire run.

Data structures

This section describes the different types of variables, arrays, and operators that are available with `awk`.

Variables

`awk` provides the ability to store the results of arithmetic and string expressions in variables for later use in the program. Referring to the previous example, consider printing the population density for each country in the file countries:

```
{ print $1, (1000000 * $3)/($2 * 1000) }
```

(Recall that in this file the population is in millions and the area is in thousands of square miles.) The result provides the number of people per square mile:

```
Russia 30.289
Canada 6.23053
China 234.561
USA 60.5809
Brazil 35.3013
Australia 4.71698
India 501.97
Argentina 24.2537
Sudan 19.6281
Algeria 19.5652
```

To improve the formatting, use `printf` as follows:

```
{printf "%10s %6.1f", $1,
 (1000000 * $3)/($2 * 1000) }
```

produces

```
Russia       30.3
Canada       6.2
China        234.6
USA          60.6
Brazil       35.3
```

```
Australia    4.7
India        502.0
Argentina    24.3
Sudan        19.6
Algeria      19.6
```

`awk` performs arithmetic internally in floating point. The arithmetic operators are `+`, `-`, `*`, `/`, and `%` (modulo or remainder).

To compute the total population and number of countries from `Asia`, you can write

```
/Asia/ { pop = pop + $3; n = n + 1 }
END { print "total population of", n,
 "Asian countries is", pop }
```

which produces

```
total population of 3 Asian countries is 1765
```

Besides writing

```
{ pop = pop + $3; n = n + 1 }
```

you can write

```
{ pop += $3; ++n }
```

The operators `++`, `--`, `-=`, `/=`, `*=`, `+=`, and `%=` function the same in `awk` as the corresponding operations in C. The statement

```
x += y
```

has the same effect as

```
x = x + y
```

but `+=` is shorter and runs slightly faster. The same is true of the `++` operator; it adds one to the value of a variable. The increment and decrement operators `++` and `--` (as in C) can be used as prefix or as postfix operators. These operators are also used in expressions.

Initialization of variables

In the previous example, neither `pop` nor `n` was initialized, yet everything worked properly. This is because (by default) variables are initialized to the null string, which has a numeric value of 0. This eliminates the need for most initialization of variables in `BEGIN` sections.

You can use references to variables in the pattern, as in

```
maxpop < $3 {
     maxpop = $3
     country = $1
 }
END { print country, maxpop }
```

which finds the country with the largest population:

```
China 866
```

Assignment operators

As described in the preceding section, "Initialization of Variables," variables can be created by virtue of an assignment operation. Other assignment operators are also available besides equal (=). Like equal, these assignment operators can be used to store a value into a variable or an element of an array (see the section that follows for more about arrays). The operators are `++`, `--`, `-=`, `/=`, `*=`, `+=`, and `%=`. They perform the same function in `awk` as the corresponding operations in C. The statement

```
x += y
```

has the same effect as

```
x = x + y
```

Most of the assignment operators are binary operators that require a variable name, followed by the operator and the value or value-producing expression, as follows:

variable-name assign-op expression

However, the `++` and `--` assignment operators are unary operators used to increment or decrement the value that was previously stored in the variable. The format they take is either a prefix or postfix format (with no space between the operator and the variable name):

unary-assign-op variable-name
variable-name unary-assign-op

So you can change

```
{ pop = pop + $3; n = n + 1 }
```

to the following line

```
{ pop += $3; ++n }
```

which uses the prefix increment operator to obtain the same processing.

Arrays

`awk` provides one-dimensional arrays as well as ordinary variables, although a name can not be both a variable and an array.

Array elements are not declared; they spring into existence when the program first encounters them. Subscripts can have any non-null value, including alphanumeric strings. As an example of a conventional numeric subscript, the statement

```
x[NR] = $0
```

assigns the current input line to the `NR`th element of the array `x`. In fact, it is possible in principle (though perhaps slow) to process the entire input in arbitrary order with the following `awk` program:

```
 { x[NR] = $0 }
END { action }
```

The first line of this program reads each input line into the array `x`.

When run on the file countries, the program

```
{ x[NR] = $1 }
```

produces an array of elements with

```
x[1] = "Russia"
x[2] = "Canada"
x[3] = "China"
```

and so forth. Arrays can also be indexed by non-numeric values, thus giving `awk` a capability rather like the associative memory of Snobol tables. For example, you can write

```
/Asia/ { pop["Asia"] += $3 }
/Africa/ { pop["Africa"] += $3 }
END { print "Asia=" pop["Asia"],
 "Africa=" pop["Africa"] }
```

which produces

```
Asia=1765 Africa=37
```

Notice the concatenation. Also, any expression can be used as a subscript in an array reference. Thus,

```
area[$1] = $2
```

uses the first field of a line (as a string) to index the array area.

Although `awk` does not support multidimensional arrays as such, you can simulate them using one-dimensional arrays. For example,

```
for ( i = 1; i <= 10; i++)
for ( j = 1; j <= 10; j++)
mult[i,j] = .\|.\|.
```

creates an array whose subscripts have the apparent form `i,j` (that is, `1,1`; `1,2`; and so forth) and thus simulates a two-dimensional array.

Internally, these subscript strings are transformed so that the comma separator is converted into the value of the variable `SUBSEP`, which, by default, is the ASCII character for code 28. Since this character does not normally appear in input text, variables can be used as subscripts, even when such variables are assigned comma-containing strings.

A special form of the `for` loop is available to iterate once through the elements of an array. You can even use it along with `SUBSEP` to reproduce the original array subscripts used with each element assignment through the following program:

```
BEGIN {
 array["one"] = "1"
 array["two,three","four"] = "2comma3, 4"
 for ( j in array ) {
      split(j,x,SUBSEP)
      printf "array["
      sep = ""
      for ( k in x ) {
           printf sep "\"" k "\""
                 sep = ","
            }
       printf "]\n"
       }
exit
}
```

This program produces the following output:

```
array["2","1"]
array["1"]
```

Built-in variables and arrays

The following list shows all of the variables maintained by `awk`:

`ARGC`	Number of command-line arguments.
`ARGV`	Array containing `ARGC` elements, one for each of the arguments that appeared on the `awk` command line.
`FILENAME`	The name of the input file currently being read. This is useful because `awk` commands can accept multiple input files, as in `awk` [*flag-options*]... *file1 file2 file3*
`FNR`	Input record number counting from the first line of the current input file.

`FS`	Input field separator; by default it is set to a blank or a tab.
`NF`	Number of fields in the current record.
`NR`	Number of command-line arguments.
`OFS`	Output-field separator; by default it is set to a blank.
`ORS`	Output-record separator; by default it is set to the newline character.
`OFMT`	The format for printing numbers; with the print statement, by default it is `%.6g`.
`RLENGTH`	Length of the string matched through the use of the match function.
`RSTART`	Beginning position of string matched through the use of the match function.
`RS`	Input record separator; by default it is set to the newline character.
`SUBSEP`	Separator for array subscripts.
`$0`	The current input record complete with unstripped field separators.
`$digit`	These variables reference fields in the current input record where `$1` contains field one, `$2` contains field two, and so on up through to the final field parsed. If only one field is found, the value of `$1` is the same as `$0`.

Expressions

Expressions are allowable wherever `awk` normally expects a value. Conditionals (or conditional expressions) can be used wherever a true or false value is expected. For more information about conditionals, see "Components of `awk` Programs," earlier in this chapter.

Arithmetic expressions can occur wherever `awk` expects a number value. Likewise, string values can occur wherever `awk` expects a string value. In cases where you supply an expression that evaluates to a value of the wrong type for a given context, the result is automatically converted into the appropriate data type as described in "Determination of Data Type" later in this section.

To manipulate numeric values, the arithmetic operators can be used. `awk` performs arithmetic internally in floating point. The operators are outlined in Table 9-1.

Table 9-1 Arithmetic operators

Symbol	Description
+	Unary and binary plus
-	Unary and binary minus
*	Multiplication
/	Division
%	Modulus
(...)	Grouping
x^y	Exponential operator

Variable assignments can be requested along with each of these arithmetic operations by using the hybrid operators for both arithmetic and assignment (see "Assignment Operators" earlier in this chapter). Table 9-2 outlines these operators.

Table 9-2 Assignment operators

Symbol	Description
=	Assign right side value to left side
+=	Increment left side by value of right side
-=	Decrement left side by value of right side
*=	Multiply left side by value of right side
=	Divide left side by value of right side
%=	Take modulus of left side by value of right side
++	Increment operand by one before/after taking current value
--	Decrement operand by one before/after taking current value

Conditional expressions (also known as relational expressions) can be used to generate true or false results. These operators are shown in Table 9-3.

Table 9-3 Relational operators

Symbol	Description
<	Less than
<=	Less than or equal
==	Equal
!=	Not equal
>	Greater than
>=	Greater than or equal

All but two of the operators listed thus far are binary operators, requiring a left-side component and a right-side component:

left-component binary-operator right-component

The left component, the right component, or both, can be replaced by other expressions, as follows:

left-component	*binary-operator*	*binary-expression*
expression	*binary-operator*	*right-component*
expression	*binary-operator*	*expression*

The left or right components can be references to variable names, number or string literals, calls to functions, other subexpressions, or any combination of these. Parentheses can be used to establish the order of evaluation when the default operator precedence is not desired.

The `++` (increment) and `--` (decrement) operators are unary operators, requiring either a left component or a right component, but not both:

left-component	*unary-op*
unary-op	*right-component*

(Note that there is no space between these unary operations and the component they affect, yielding `count++` and `++count` to increment the variable `count` either after or before use.)

The left component or right component is often a variable name, but it can also be a number literal, a function call, or any subexpression enclosed in parentheses.

By nesting expressions as the right component or left component of either binary or unary operator expressions, computations can be created to any level of complexity.

You can also use logical operators and pattern-matching operators for regular expressions, as shown in Tables 9-4 and 9-5.

Table 9-4 Logical operators

Symbol	Description
`||`	OR
`&&`	AND
`!`	NOT
`?:`	If-then-else contruct. For example, `x?y:z` yields `y` if `x` is true, else yields `z`.

Table 9-5 Regular expression pattern-matching operators

Symbol	Description
`~`	Matches
`!~`	Does not match

The nesting of expressions is also possible for expressions using the relational operators, with an added restriction: the operator at the uppermost level should be the AND (`&&`) or OR (`||`) binary operators or the unary NOT (`!`) operator. For more detailed information, see the next section, "Combining True-or-False Expressions."

If you create an expression without any operators, the expression must be a reference to a variable or array, a literal string or number value, or a function. See the sections "Variables" and "Arrays" earlier in this chapter for more information. See "Numeric Constants" and "String Constants" later in this chapter for more information about literal values. See "Built-in String Functions," "Built-in Numeric Functions," and "User-Defined Functions" later in this chapter for information regarding functions.

Combining true-or-false expressions

Whether you are specifying the contents of a pattern or an action, expressions that return true or false can be combined using the logical operations `||` (OR), `&&` (AND), `!` (NOT), and parentheses. Using the countries example, the program

```
$2 >= 3000 && $3 >= 100
```

selects lines where both area and population are large:

```
Russia 8650 262 Asia
China 3692 866 Asia
USA 3615 219 N. America
Brazil 3286 116 S. America
```

The program

```
$4 == "Asia" || $4 == "Africa"
```

selects lines with `Asia` or `Africa` as the fourth field. An alternate way to write this last expression is with a regular expression:

```
$4 ~/Asia|Africa/
```

The operators `&&` and `||` guarantee that their operands are evaluated from left to right; evaluation stops as soon as truth or falsehood is determined.

Implied concatenation operations

Although none of the lists showing `awk` operation symbols includes a symbol that represents string concatenation, this operation is nevertheless invoked regularly within expressions.

When separated by a space, string or number expressions are concatenated into one string with no intervening spaces. If number expressions are used, they are evaluated arithmetically and converted into strings before concatenation is performed. The second of the three following action statements performs a concatenation:

```
{
 x = "hello"
 x = x ", world"
 print x
}
```

This prints the usual:

```
hello, world
```

With input from the file countries, the program

```
/^A/ { s = s $1 " " }
END { print s }
```

prints

```
Australia Argentina Algeria
```

Determination of data type

Variables (and fields) take on numeric or string values according to context. For example, in

```
pop += $3
```

`pop` is presumably a number, while in

```
country = $1
```

`country` is a string. In

```
maxpop < $3
```

the type of `maxpop` depends on the data found in `$3`. It is determined when the program is run. In general, each variable and field is potentially a string or a number or both at any time. When a variable is set by the assignment

`var =` *expression*

its type is set to that of *expression*. (Assignment also includes `+=`, `++`, `-=`, and so forth.) An arithmetic expression is of the type number; a concatenation of strings is of the type string.

In comparisons, if both operands are numeric, `awk` makes the comparison numerically. Otherwise, operands are coerced to strings, if necessary, and the comparison is made on strings.

The type of any expression can be coerced to numeric by maneuvers such as

expression `+ 0`

and to string by

expression `""`

This last expression is a string concatenated with the null string. If a string cannot be converted to a number without errors, `awk` converts it to zero.

Built-in string functions

The `length` function computes the length of a string of characters and the usage format is as follows:

`length(`*string*`)`

For example, with input taken from the file `countries`, the following `awk` program prints the longest country name:

```
length($1) > max { max=length($1); name=$1 }
END { print name }
```

If you don't include a parenthetical argument, `length` returns the length of the current input record. The following program prints each record preceded by its length:

```
{ print length, $0 }
```

In this case `length` is equivalent to `length($0)`.

The function

`split(`*string*`, `*array* `[,`*sep*`] )`

assigns the fields of *string* to successive elements of the array *array*. When *sep* is missing, the separator used is that given by the built-in variable `FS`. For example:

```
split("Now is the time", w)
```

assigns the value `Now` to `w[1]`, `is` to `w[2]`, `the` to `w[3]`, and `time` to `w[4]`. All other elements of the array `w`, if any, are set to the null string.

When `awk` evaluates a `split` function, it returns the number of array elements created. Accordingly,

`count = split(`*string*`, `*array*`, `*sep*`)`

assigns the number of elements initialized in array to the variable `count`. Use assignments of this form when you must know how many elements a string is split into.

When *sep* is present, it must be a single character enclosed in double quotation marks but only its first character is used as the field separator. For instance, if you use the following three lines (`\t` is the tab character),

```
{split("Now is+the time", w, "+")}
{split("This~is~not~the~end", x , "~")}
{print w[1],x[3],w[2] }
```

the output is

```
Now is not the time
```

The `substring` function

`substr(`*string*, *position*, *length*`)`

produces the substring of *string* that begins at column *position* and is, at most, *length* characters long. If the *length* is omitted, the returned substring extends to the end of *string*. For example, you can abbreviate the country names in the file `countries` by running the `awk` program

```
{ $1 = substr($1, 1, 3); print }
```

which produces

```
Rus 8650 262 Asia
Can 3852 24 N. America
Chi 3692 866 Asia
USA 3615 219 N. America
Bra 3286 116 S. America
Aus 2968 14 Australia
Ind 1269 637 Asia
Arg 1072 26 S. America
Sud 968 19 Africa
Alg 920 18 Africa
```

If string is a number, `substr` uses its string representation; for example, `substr(123456789,3,4)` is `3456`.

The function

`index(`*string*, *lookup-string*`)`

returns the left-most position where *lookup-string* occurs inside *string*, or zero if *string* does not contain *lookup-string*.

A variant on the `index` function is `match`, the format of which is

`match(`*string*`,` *pattern*`)`

which returns the left-most position where a substring of *string* is matched by the regular expression *pattern*, or zero if no match is found.

The functions `gsub` and `sub`

`[g]sub(`*pattern*`,` *new-string*`,` *string*`)`

replace occurrences of substrings within *string* that are sought by the regular expression *pattern* with *new-string*. To replace only the first substring sought by the regular expression, use `sub`. To replace all nonoverlapping substrings sought by the regular expression, use the global-substitute function `gsub`.

The function

`sprintf(`*format-string*`,` *expr* `[,` *expr*`]...)`

formats expressions as the `printf` statement does, but assigns the resulting expression to a variable instead of sending the results to the standard output. For example:

```
x = sprintf("%10s %6d ", $1, $2)
```

sets `x` to the string produced by formatting the values of `$1` and `$2`. The `x` can then be used in subsequent computations.

Built-in numeric functions

`awk` also provides the following mathematical functions:

`atan2(`*number*`)`

`cos(`*radians*`)`

`exp(`*number*`)`

`int(`*number*`)`

`log(`*number*`)`

`rand()`

`srand(`*seed-number*`)`

`sin(`*radians*`)`

`sqrt(`*number*`)`

For the most part, these functions are the same as those of the C library, returning the same errors as those in `libc`. (See "C Special Libraries" in *A/UX Programming Languages and Tools*, Volume 1.) The result returned by the random number function is a value greater than 0 and less than or equal to 1. The `int` corresponds to the C library floor function because of the way it handles negative numbers.

Lexical conventions

All `awk` programs are made up of lexical units called tokens. `awk` uses eight types of tokens:

- numeric constants
- string constants
- keywords and built-in variables
- identifiers
- operators
- record and field tokens
- comments (discussed previously)
- separators

Precise specifications of each token are given in the following sections.

Numeric constants

A numeric constant is either a decimal constant or a floating constant. A decimal constant is a non-null sequence of digits containing, at most, one decimal point, as in

```
12
12.
1.2
.12
```

A floating constant is a decimal constant followed by `e` or `E` followed by an optional `+` or `-` sign followed by a non-null sequence of digits, as in

```
12e3
1.2e3
1.2e-3
1.2E+3
```

String constants

A string constant is a sequence of zero or more characters surrounded by double quotation marks, as in

```
"armadillo"
"a"
"ab"
"12"
```

A double quotation mark can be put into a string by preceding it with the backslash (\), as in

```
"He said, \"Sit!\""
```

A newline is put in a string by using `\n` in its place. No other characters need to be escaped except `\` itself. Strings can be (almost) any length.

Predefined variables, reserved keywords, and reserved function names

Table 9-6 lists certain character strings that have special meaning to `awk`. There are three types of these character strings:

1. **Predefined variables** are variables defined by `awk` that have special meanings. The meaning of these variables is explained in "Special Variables."
2. **Reserved keywords** are a special set of character strings used in `awk` statements. Reserved keywords cannot be used as variables.

3. **Reserved function names** are a special set of character strings used to invoke built-in `awk` functions. These functions are discussed in "Built-in Functions" earlier in this chapter.

Table 9-6 Reserved strings

Predefined variables	Reserved keywords	Reserved function name
BEGIN	break	exp
END	close	getline
FILENAME	continue	index
FS	exit	int
NF	for	length
NR	in	flog
OFS	next	split
ORS	number	sprintf
OFMT	print	sqrt
RS	printf	substr
$0	string	
$i	while	

Identifiers

Identifiers in `awk` serve to denote variables and arrays. An identifier is a sequence of letters, digits, and underscores beginning with a letter or an underscore. Uppercase and lowercase letters are different.

Record and field tokens

`$0` is a special variable whose value is the current input record. `$1`, `$2`, and so on, are special variables whose values are the first field, the second field, and so on, respectively, of the current input record. The keyword `NF` (number of fields) is a special variable whose value is the number of fields in the current input record. Thus, `$NF` has as its

value the value of the last field of the current input record. Notice that the first field of each record is numbered 1 and that the number of fields can vary from record to record. None of these variables is defined in the action associated with a `BEGIN` or `END` pattern, where there is no current input record.

The keyword `NR` (number of records) is a variable whose value is the number of input records read so far. The first input record read is 1. At `END` it contains the total number of input lines.

Separators

The `awk` language provides two data-separator variables to assist in parsing information, the record-separator and field-separator variables.

Record separators

The keyword `RS` (record separator) is a variable whose value is the current record separator. The value of `RS` is initially set to newline, indicating that adjacent input records are separated by a newline. Keyword `RS` is changed to any character *c* by including the assignment statement

```
RS = "c"
```

in an action.

Field separator

The keyword `FS` (field separator) is a variable indicating the current field separator. Initially, the value of `FS` is a blank, indicating that fields are separated by white space—that is, any sequence of blanks and tabs. Keyword `FS` can be changed to any single character *c* by including the assignment statement

```
FS = "c"
```

in an action or by using the flag option `-Fc`. Two values of *c*, space and `\t`, have special meaning. The assignment statement

```
FS = " "
```

makes white space (blank spaces or tabs) the field separator; on the command line, `-F"\t"` makes a tab the field separator.

If the field separator is not a blank, there is a field in the record on each side of the separator. For instance, if the field separator is 1, the record `1XXX1` has three fields. The first and last are null, and the value of the second is `XXX`. If the field separator is blank, fields are separated by white space, and none of the `NF` fields are null; that is, record `1XXX1` has one field, not three, as in the previous case.

Multiline records

The assignment

```
RS = ""
```

as part of the action associated with a `BEGIN` pattern makes an empty line the record separator. It also makes a sequence of blanks, tabs, and possibly a newline, the field separator. With this setting, none of the first fields of any record is null, as discussed earlier.

Output record and field separators

The value of `OFS` (output field separator) is the character or string separating output fields. It is put between fields by print. The value of `ORS` (output record separator) is put after each record by print. Initially, `ORS` is set to a newline and `OFS` to a space. These values can be changed to any string by assignments such as the following two:

```
ORS = "abc"
OFS = "xyz"
```

Separators and braces

Tokens in `awk` are usually separated by non-null sequences of blanks, tabs, and newlines, or by other punctuation symbols, such as commas and semicolons. Braces (`{ }`) surround actions, slashes (`/ /`) surround regular expression patterns, and double quotation marks (`" "`) surround strings. Braces also can be used to group statements within actions.

Primary expressions

In `awk`, patterns and actions are made up of expressions. The basic building blocks of expressions are the following primary expressions:

- numeric constants
- string constants
- variables
- functions

Each expression has both a numeric and a string value, and defaults to one or the other, depending on context. The rules for determining the default value of an expression are explained in the following sections.

Numeric constants

A numeric constant is simply a number. The format of a numeric constant was previously defined in the section "Lexical Conventions." The value of a numeric constant is always its numeric value in decimal unless it is coerced to type string. Table 9-7 shows the result of coercing various numeric constants to type string. Coercion of a numeric constant can occur explicitly as defined in "Type" or implicitly within the context of an expression.

Table 9-7 Values for sample numeric constants

Numeric constant	Numeric value	String value
`0`	`0`	`0`
`1`	`1`	`1`
`.5`	`0.5`	`5`
`.5e2`	`50`	`50`

String constants

A string constant is simply a series of characters enclosed in double quotation marks. The format of a string constant was defined in "Lexical Conventions" earlier in this chapter.

The value of a string constant is the contents of the string itself unless it has been coerced to type numeric. The numeric value of a string coerced to type numeric depends on the contents of the string: If the string is composed entirely of numbers (either decimal or floating-point format), its numeric value is the number contained in the string. If the string does not contain a recognizable decimal or floating-point number, its numeric value is zero. Table 9-8 shows the result of coercing various string constants to type numeric. Coercion of a string constant can occur explicitly as defined in "Type" or implicitly within the context of an expression.

Table 9-8 Values for sample string constants

String constant	Numeric value	String value
`""`	0	*null string*
`" "`	0	*space*
`"a"`	0	`a`
`"XYZ"`	0	`XYZ`
`"0"`	0	`0`
`"1"`	1	`1`
`".5"`	0.5	`.5`
`".5e2"`	50	`.5e2`

Variables

A **variable** or **var** is in one of the following forms:

identifier

identifier`[` *expression* `]`

`$`*term*

The numeric value of any uninitialized variable is `0`, and the string value is the empty string. An *identifier* by itself is a simple variable. A variable of the form

identifier`[` *expression* `]`

represents an element of an associative array named by *identifier*.

The string value of *expression* is used as the index into the array. The default value of *identifier* or *identifier[expression]* is determined by context.

The variable `$0` refers to the current input record. Its string and numeric values are those of the current input record. If the current input record represents a number, the numeric value of `$0` is the number and the string value is the literal string. The default value of `$0` is string unless the current input record is a number. `$0` cannot be changed by assignment.

The variables `$1` and `$2` refer to fields 1 and 2 of the current input record. The string and numeric values of `$i` for `1<=`*i*`<=NF` are those of the *i*th field of the current input record. As with `$0`, if the *i*th field represents a number, the numeric value of `$`*i* is the number and the string value is the literal string. The default value of `$`*i* is a string unless the *i*th field is a number. The `$`*i* can be changed by assignment. The value of `$0` is then changed accordingly, but the results might not be apparent unless `NF` is changed to at least *i*.

In general, `$`*term* refers to the input record if term has the numeric value `0` and to field *i* if the greatest integer in the numeric value of term is *i*. If *i*<0 or if *i*>=100, then accessing `$`*i* causes `awk` to produce an error diagnostic. If `NF`<*i*<=100, then `$`*i* behaves like an uninitialized variable. Accessing `$`*i* for *i* > `NF` does not change the value of `NF`.

Functions

The `awk` language has a number of built- in functions that perform common arithmetic and string operations.

`exp` `[(`*expression*`)]`

`int` `[(`*expression*`)]`

`log` `[(`*expression*`)]`

`sqrt` `[(`*expression*`)]`

These functions (`exp`, `int`, `log`, and `sqrt`) compute the exponential, integer part, natural logarithm, and square root, respectively, of the numeric value of expression. The (expression) can be omitted; then the function is applied to `$0`. The default value of an arithmetic function is numeric.

`getline`

`index (`*expression1*`,` *expression2*`)`

`length [(`*expression*`)]`

`split (`*expression*`,` *identifier* `[, "`*separator*`"])`

`sprintf [("`*format*`",` *expression1* `[,`*expression2* `...] )]`

`substr (`*expression1*`,` *expression2* `[,` *expression3* `] )`

These functions (`getline`, `index`, `length`, `split`, `sprintf`, and `substr`) perform string operations. See "Built-in String Functions" earlier in this chapter for more details.

Terms

Various arithmetic operators are applied to primary expressions to produce larger syntactic units called **terms.** All arithmetic is done in floating point. A term has one of the following forms:

primary expression
term1 binop term2
unop term
incremented var
(term)

Binary terms

In a term of the form

term1 binop term2

binop can be one of the five binary arithmetic operators `+` (addition), `-` (subtraction), `*` (multiplication), `/` (division), or `%` (modulus). The binary operator is applied to the

numeric value of the operands *term1* and *term2*, and the result is the usual numeric value. This numeric value is the default value, but it can be interpreted as a string value (see "Numeric Constants" earlier in this chapter). The operators `*`, `/`, and `%` have higher precedence than `+` and `-`. All operators are left associative.

Unary terms

In a term of the form

unop term

unop can be unary `+` or `-`. The unary operator is applied to the numeric value of *term*, and the resulting numeric value is the default value. However, it can be interpreted as a string value. Unary `+` and `-` have higher precedence than `*`, `/`, and `%`.

Incremented variables

An incremented variable has one of the following forms:

`++`*var*

`--`*var*

var`++`

var`--`

That is, it can be either *pre-* or *post-incremented.*

The form `++`*var* has the effect of the assignment

var = *var* `+ 1`

and so has the value *var*`+1` before it is further evaluated or assigned. Similarly, the form `--`*var* has the effect of the assignment

var = *var* `- 1`

and so has the value *var*`-1` before it is further evaluated or assigned.

The form *var*`++` has the same value as *var* before it is evaluated or assigned, and after that it has the effect of the assignment

var = *var* `+ 1`

Similarly, the form *var*`--` has the same value as *var* before it is evaluated or assigned, and after that it has the effect of the assignment

var = *var* `- 1`

The default value of an incremented *var* is numeric. You shouldn't use the `++` or `--` operators where the incremented variable is used more than once (such as `a = b++ * b`), since the results are indeterminate.

Terms with parentheses

Parentheses are used to group terms in the usual manner.

Expressions

An `awk` expression is one of the following:

term

term1 term2 ...

var asgnop expression

Concatenation of terms

In an expression of the form *term1 term2*, the string values of the terms are concatenated. If the terms are numeric expressions, they are first evaluated and then also treated as strings; that is, the default value of the resulting expression is a string value that can be interpreted as a numeric value. Concatenation of terms has lower precedence than binary `+` and `-`. For example, the expression

```
1+2 3+4
```

has the string (and numeric) value 37.

Assignment expressions

An assignment expression is one of the form

var asgnop expression

where *asgnop* is one of the six assignment operators (`=`, `+=`, `-=`, `*=`, `/=`, `%=`, `++`, `--`) (see "Operators" earlier in this chapter).
The default value of *var* is the same as that of expression.

In an expression of the form

var = *expression*

the numeric and string values of *var* become those of expression.

An expression of the form

var op = *expression*

is equivalent to

var = *var op expression*

where *op* is one of the arithmetic operators (see "Operators" earlier in this chapter).

The *asgnops* are right associative and have the lowest precedence of any operator. Thus, the assignment

```
a += b *= c - 2
```

is interpreted as

```
a += ( b *= ( c - 2 ) )
```

which is equivalent to the sequence of assignments

```
b = b * (c - 2)
a = a + b
```

Part 3 Math Tools

A/UX provides two specialized tools for handling arbitrary precision arithmetic, `dc` and `bc`. The `dc` program is an interactive desk calculator program. It has provisions for manipulating scaled fixed-point numbers and for input and output in bases other than decimal. `bc` is a specialized language and compiler for handling arbitrary precision arithmetic using the `dc` program. The following two chapters describe these tools.

10 dc: A Desk Calculator

Using dc / 10-2
Programming dc / 10-9

dc is an interactive desk calculator program for handling arbitrary-precision integer arithmetic. It has provisions for manipulating scaled fixed-point numbers and for input and output in bases other than decimal.

The dc program works like a stacking calculator using reverse Polish notation. Ordinarily, dc operates on decimal integers; however, the input base, output base, and scale can be set according to user specifications. Because dc is based on a dynamic storage allocator, number size is limited only by available core storage.

dc can also be used in conjunction with bc, a high-level language and compiler designed specifically as a front end for dc. Complex functions can be defined and saved in a file for later execution through bc. When a program is executed, bc compiles the input and automatically pipes it to the dc interpreter, which produces the final result. See the next chapter, "bc: A Basic Calculator," in this manual for more information.

Using dc

To begin using dc, simply type its name to the shell:

```
dc
```

Anything you then enter is interpreted as dc input, up to an end-of-file (CONTROL-D). You also can exit from dc by using the q command, discussed later.

For very complex computations, you might find it more efficient to place the instructions into a file. You can then pass the filename as an argument to the dc command:

dc *filename*

dc reads and executes the contents of the *filename* argument before accepting further commands from the keyboard.

dc operates like a stacking calculator using reverse Polish notation. Initially, the value of a number is pushed onto the stack. The top two values on the stack can then be added (+), subtracted (-), multiplied (*), divided (/), remaindered (%), or exponentiated (^), according to the current operator. The two entries are popped off the stack, and the result is pushed on the stack in their place.

Similarly, the top value on the stack can be duplicated, removed, stored in a register, and so forth. For the full list of operations, see the following section.

Command syntax

You can have any number of commands on a line. Blanks and newline characters are ignored, except when used to delineate numbers and in places where a register name is expected. Tabs are not allowed.

A number is an unbroken string of digits 0 through 9 and uppercase letters A through F (treated as digits with values 10 through 15, respectively). A negative number can be indicated by preceding a number with an underscore (_). Numbers also can contain decimal points.

To perform simple operations, you can use the following format:

```
24.2 56.2 + p
```

The `p` command instructs `dc` to print the result of the computation (in this case, an addition). Here is an example of a more complex problem, using a variety of commands:

```
[la1+dsa*pla10>y]sy
0sa
lyx
```

This example prints the first ten values of the factorial function (that is, 1! through 10!). To fully understand how it does so, please see "Programming `dc`" later in this chapter.

Operators

Table 10-1 shows the operators that can be used in `dc` expressions:

Table 10-1 `dc` operators

Operator	Function
^	Exponentiation
*	Multiplication
%	Remaindering modulus (integer result truncated toward zero)
/	Division
+	Addition
-	Subtraction
v	Square root

Relational operators

`dc` allows the following relational operators (also referred to as testing commands):

```
<x >x =x !<x !>x !=x
```

These cause the top two elements of the stack to be popped and compared. Register `x` is executed if the top two elements of the stack satisfy the stated relation. The exclamation point indicates negation.

`dc` command set

The following sections describe the `dc` commands in detail, categorized by subject. At the end of the categorized sections is a quick-reference list of all `dc` commands, with brief descriptions of each.

Input/output format and base

The input and output bases affect only the interpretation of numbers on input and output. They have no effect on internal arithmetic computations.

Large numbers are generated with 70 characters per line; a backslash (`\`) indicates a continued line. All choices of input and output bases work correctly, although not all are useful. A particularly useful output base is 100000, which has the effect of grouping digits in fives. Bases of 8 and 16 are used for decimal-octal or decimal-hexadecimal conversions.

Input conversion and base

Numbers are converted to their internal representation as they are read in to `dc`.

`_` Negative numbers are indicated by preceding the number with an underscore (_).

`i` The `i` command can be used to change the base of the input numbers. This command pops the stack, truncates the resulting number to an integer, and uses it as the input base for all further input. The default for input base (`ibase`) is 10 (decimal) but can, for example, be changed to 8 or 16 for octal-to-decimal or hexadecimal-to-decimal conversions.

`I` The `I` command pushes the value of the input base on the stack.

No mechanism is provided for the input of arbitrary numbers in bases less than 1 or greater than 16. The hexadecimal digits A through F correspond to the numbers 10 through 15, regardless of input base.

Output commands

`p` The `p` command causes the top of the stack to be printed. It does not remove the top of the stack.

`f` The `f` command prints the contents of all of the stack registers.

`o` The `o` command is used to change the output base (`obase`). This command uses the top of the stack truncated to an integer as the base for all further output. The default output base is 10 (decimal).

`O` The `O` command pushes the value of the output base on the stack.

Scale

`dc` can accommodate scales up to 99 decimal places. The default scale is 0.

`k` The `k` command sets the scale to the number on the top of the stack, truncated to an integer.

`K` The `K` command can be used to push the value of `scale` on the stack. The value of `scale` must be greater than or equal to 0 and less than 100.

The rules governing how the scale of a result is resolved for the different operations are as follows:

Operator	*Scale*
`^`	The scale of the result is the sum of the scales of the two operands. If this exceeds the value of `scale`, it is truncated to that value.
`*`	The scale of the result is the sum of the scales of the two operands. If this exceeds the value of `scale`, it is truncated to that value.
`%`	The scale of the remainder is the maximum of the dividend scale and quotient scale, plus the divisor scale.
`/`	The scale of the result is the value of `scale`. You must specify a `scale` value for any scale to occur.
`+`	The scale of the result is the larger scale of the two operands.
`-`	The scale of the result is the smaller scale of the two operands.
`v`	The scale of the result is given the scale of the operand or the value of `scale`, whichever is larger.

Stack commands

`c` The `c` command clears the stack.

`d` The `d` command pushes a duplicate of the top number onto the stack.

`z` The `z` command pushes the stack size onto the stack.

`X` The `X` command replaces the number on the top of the stack with its scale factor.

`Z` The `Z` command replaces the top of the stack with its length.

Subroutine definitions and calls

`[]` Enclosing a string in brackets pushes the ASCII string onto the stack.

`q` The `q` command quits or (when executing a string) pops the recursion level by two.

Internal registers

Numbers or strings can be stored in internal registers or loaded on the stack from registers with the commands `s` and `l`:

`s`*x* The `s`*x* command pops the top of the stack and stores the result in register *x*. The *x* can be any character; even a blank or newline is considered a valid register name.

`l`*x* The `l`*x* command puts the contents of register *x* on the top of the stack. The *x* can be any character; even a blank or newline is considered a valid register name.

◆ **Note** The `l` command has no effect on the contents of register *x*. The `s` command, however, is destructive. ◆

Pushdown registers and arrays

◆ **Note** The following commands are intended for use by a compiler, rather than for direct use by programmers. ◆

`dc` can be thought of as having individual stacks for each register. These registers are operated on by the commands `S` and `L`:

`S`*x*	`S`*x* pushes the top value of the main stack onto the stack for the register *x*.
`L`*x*	`L`*x* pops the stack for register *x* and puts the result on the main stack.
`s` and `l`	The `s` and `l` commands also work on registers, but not as pushdown stacks. The `l` command does not affect the top of the register stack, but `s` destroys what was there before.

The commands that work on arrays are `:` and `;`.

`:`*x*	The `:`*x* command pops the stack and uses this value as an index into the array *x*. The next element on the stack is stored at this index in *x*. An index must be greater than or equal to 0 and less than 2048.
`;`*x*	The `;`*x* command loads the main stack from the array *x*. The value on the top of the stack is the index into the array *x* of the value to be loaded.

Miscellaneous commands

`!`	The `!` command interprets the rest of the line as an A/UX system command and passes it to the operating system to execute.
`Q`	The `Q` command uses the top of the stack as the number of levels of recursion to skip.

`dc` command quick reference

The following list is a quick reference to `dc` command characters and their functions:

`[...]`	Puts the bracketed character string on top of the stack.
`!`	Interprets the rest of the line as an A/UX system command. Control returns to `dc` when the command terminates.
`?`	Takes a line of input from the input source (usually the console) and executes it.
`c`	Pops all values on the stack; the stack becomes empty.
`d`	Duplicates the top value on the stack.
`f`	Prints all values on the stack and in registers.
`i` and `I`	Pops the top value on the stack and uses it as the number radix for further input. The command `I` pushes the value of the input base on the stack.
`k` and `K`	Pops the top of the stack and uses that value as a scale factor that determines the maximum number of decimal places that are maintained during multiplication, division, and exponentiation. The scale factor must be greater than or equal to zero and less than 100. The `K` command can be used to push the value of `scale` on the stack.
`l`*x* and `L`*x*	The `l` command puts the contents of register *x* on top of the stack. The initial value of a new register is treated as a zero by the command `l`, but treated as an error by the command `L`. The `L`*x* command pops the stack for register *x* and puts the result on the main stack.
`o` and `O`	The top value on the stack is popped and used as the number radix for further output. The command `O` pushes the value of the output base on the stack.
`p`	The top value on the stack is printed. The top value remains unchanged.
`q` and `Q`	Exits from the program. If executing a string, the recursion level is popped by two. If `Q` is used, the top value on the stack is popped; and the string execution level is popped by that value.
`s`*x* and `S`*x*	The top of the main stack is popped and stored in a register named *x* (where *x* can be any character). The value of register *x* is pushed onto the stack. Register *x* is not altered. `S`*x* pushes the top value of the main stack onto the stack for the register *x*.

`v`	Replaces the top element on the stack by its square root. The square root of an integer is truncated to an integer.
`x` and `X`	The `x` command assumes the top of the stack is a string of `dc` commands, removes it from the stack, and executes it. The `X` command replaces the number on the top of the stack with its scale factor.
`z` and `Z`	The value of the stack level is pushed onto the stack. The `Z` command replaces the top of the stack with its length.

Programming `dc`

By combining a few of the available constructs, such as the load, store, execute, and print commands (`l`, `s`, `x`, `p`), the `[ ]` construct to store strings, and the testing commands (relational operators), it is possible to program `dc`. For example, the following expressions instruct `dc` to print the numbers 0 through 9:

```
[lip1+sili10>a]sa
0si
lax
```

Consider the first expression in this example:

```
[lip1+sili10>a]sa
```

This first instruction makes use of the `[ ]` construct for storing strings. The entire expression is stored as a character string on top of the stack. Reading from left to right, this character array holds the following commands:

- Load the contents of register `i` on top of the stack, and print it.

  ◆ **Note** Using the print command does not remove the top of the stack. ◆

- Add (+) 1 to the value found on top of the stack, and place the result on top of the stack.
- Store the value currently found on top of the stack in register `i`.

- Load the contents of register `i` on top of the stack, then load the number 10 onto the stack. Use the testing operator `>` on these top two stack elements to see whether 10 is greater than the number that was loaded from register `i`. If 10 is greater, execute register `a`. This is the "control element" in this example, because it stops the processing of the expressions as soon as the value in register `i` is equal to 10.
- Store the character array in register `a`. The second and third lines of the example contain the expressions

  ```
  0 si
  la x
  ```

- The `0 si` instruction clears register `i` by storing 0 in that register, thereby removing any previous value it may have had.
- The `la` and `x` instructions load the contents of register `a` on top of the stack and execute it.

◆ **Note** The size of numbers in `dc` is limited only by the size of available memory. ◆

11 bc: A Basic Calculator

Using bc / 11-3
Program syntax / 11-5

bc is a specialized language and compiler for handling arbitrary-precision arithmetic. bc calls the dc calculator program to perform any actual computations. In fact, bc was designed specifically to augment dc routines for manipulating infinitely large numbers, scaled up to 99 decimal places.

Because bc is based on a dynamic storage allocator, overflow does not occur until all available core storage is exhausted. bc has a complete control structure, and can be used either in immediate mode (direct immediate input/output to and from bc) or as an interactive processor for bc programs. Consequently, complex functions can be defined and saved in a file for later execution. A small library of predefined functions is also available, among which are the sine, cosine, arctangent, logarithmic, exponential, and Bessel functions of integer order.

`bc` contains scaling provisions that permit the use of decimal-point notation, as well as input and output in bases other than base 10. Numbers can be converted from decimal to octal simply by setting the output base to 8. The limit on the number of digits that can be manipulated depends only on the amount of core storage available.

While `bc` is not intended as a complete programming language, it can be used effectively to do a number of specific tasks, most notably the following ones:

- compile large integers
- compute accurately to many decimal places
- convert numbers from one base to another base

Using bc

In this chapter, the term "bc command" refers to the command you type from the shell command line, and the term "bc program" refers to the set of calculations to be performed by the bc command. These calculations can reside in a bc program file.

bc command syntax

The bc command has the following syntax:

bc [-c] [-l] [*file*]

The -c compile-only option directs bc to output that it would normally pass as input to dc. The output is instructive but complicated.

The -l (library) option calls the set of math library functions in bc:

Function syntax	*Operation*
s(*x*)	Sine
c(*x*)	Cosine
a(*x*)	Arctangent
l(*x*)	Natural logarithm
e(*x*)	Exponential
j(*n*,*x*)	Bessel function integer order

The library option initially sets the scale (number of available decimal places after the decimal point) to 20, but this can be reset using the scale function call. See the section "scale" later in this chapter.

The *file* is an optional bc program file from which bc can read calculations.

Entering a program at the terminal

For the immediate evaluation of simple arithmetic expressions that do not involve standard `bc` library functions or any user-defined functions, simply enter the `bc` program at the terminal. For example, to perform a simple operation, first invoke `bc` and then enter the calculation to be performed:

```
bc
142857 + 285714
```

`bc` then responds immediately with the result

```
428571
```

Program files

For more complicated calculations, you might find it more efficient to define the functions or procedures in a program file. You would then pass the filename as an argument to the `bc` command:

`bc` *filename*

`bc` then reads and executes the contents of the named file before accepting further commands from the keyboard.

Exiting from `bc`

To exit from `bc`, even when using a command file, you must issue a `quit` or an end-of-file character (see `stty`(1) in *A/UX Command Reference* for more information). Unless you use the syntax `bc <` *filename*, `bc` does not quit when it reaches the end of the program file. If no `quit` statement is given, `bc` simply waits for further instructions, and your shell prompt is not returned.

To exit, you can either place a `quit` statement at the end of your file or enter `quit` or your end-of-file character directly when `bc` completes the file. Your end-of-file character can still be used as an interrupt and terminate signal while the file is being processed.

The `quit` statement is not treated as an executable statement, and so cannot be used in a function definition or in an `if`, `for`, or `while` statement.

Program syntax

The syntax of a `bc` program is very similar to that of a C-language program. In general, statements and control structures are identical in `bc` and in C. A good example of this similarity is the manner in which a `bc` function is defined. The following program defines a function that computes the approximate value of the exponential function and prints the result for the first ten integers. The pieces of this example are discussed in individual sections that follow.

```
scale = 10
define e(x) {
     auto a,b,c,i,s
     a = 1
     b = 1
     s = 1
     for(i=1; 1==1; i++) {
          a = a*x
          b = b*i
          c = a/b
          if (c == 0) return(s)
          s = s+c
        }
   }
for(i=1; i<=10; i++) e(i)
```

Comments

The characters `/` and `*` introduce a comment that terminates with the characters `*` and `/`. Anything between the asterisks is ignored by the `bc` compiler.

Constants

Constants are primitive expressions and consist of arbitrarily long numbers with an optional decimal point. The hexadecimal digits A through F are also recognized as digits with values 10 through 15, respectively.

Keywords

The following terms are reserved as `bc` keywords, and cannot be used other than for their predefined purposes:

`auto`	`for`	`length`	`return`	`while`
`break`	`ibase`	`obase`	`scale`	
`define`	`if`	`quit`	`sqrt`	

Identifiers

In `bc`, an identifier is a character, or sequence of characters, that names an expression. The identifier is the "place" where the value of that expression is stored. Therefore, identifiers are legal on the left side of an assignment statement.

`bc` has three kinds of identifiers:

- simple identifiers
- function calls
- array, or subscripted, variables

All three types should be indicated with single lowercase letters. Identifier names do not conflict; a `bc` program can have a simple variable identifier named `x`, an array named `x`, and a function named `x`, all of which are separate and distinct.

Defining functions

Functions are specified by a single lowercase letter, followed immediately by a set of parentheses:

```
a()
```

Since function names are permitted to coincide with simple variable names, the parentheses indicate the difference between a function and a variable, and provide a means of passing arguments to the function. Twenty-six different defined functions are permitted in addition to the 26 variable names.

A function is defined in the following manner:

```
define a(x) {
        defining statements
        return
    }
```

The word `define` initiates the function definition; `a(x)` names the function and indicates that the function requires one argument; the left brace opens the body of the definition and must occur on the same line as the `define` keyword; `return` returns control to the calling function; and the right brace closes the definition. The body of the definition must contain one or more statements, and must begin and close with a left and right brace, respectively.

Function calls and function arguments

A function call consists of the function name followed by parentheses, which in turn should contain any required arguments to be passed to the function. Individual arguments should each be separated by commas. Functions with no arguments are called and defined using empty parentheses. If a function is called with the wrong number of arguments, the result is unpredictable.

All function arguments are passed by value, and as a result the values remain discrete, local to the called function. Therefore, changes made to the argument values within the called function do not alter the original parameters outside the function.

The `return` statement

Return of control from a function occurs when a `return` statement is executed, or when the end of the function is reached. The `return` statement can take either of the following two forms:

```
return
return(x)
```

In the first case, the value returned from the function is 0; in the second, the value returned from the function is the expression in parentheses.

Automatic variables

Automatic variables are allocated space and initialized to zero on entry to the function, and thrown away on return (exit). The values of any similarly named variables outside the function are not disturbed. Functions can be called recursively and the automatic variables at each level of call are protected.

It should be noted, however, that automatic variables in `bc` do not work exactly the same way as they do in the C language. On entry to a function, the old values of automatic variables or parameters named previously are pushed onto a stack. Until return is made from the function, reference to these names refers only to the new values.

Variables used in a function can be declared as automatic by a statement of the form

```
auto x,y,z
```

There can be only one such `auto` statement in a function, and it must be the first statement in the definition.

The following example is a function definition that uses an automatic variable:

```
define a(x,y) {
        auto z
        z = x*y
        return(z)
    }
```

When called, the value of this function `a` is the product of its two arguments, `x` and `y`. Consequently, the input

```
a(7,3.14)
```

sends the result, 21.98, to the standard output. Using this same function, the input

```
z = a(a(3,4),5)
```

sends the result, 60, to the standard output.

Global variables

There are only two storage classes in `bc`: automatic variables and global variables. Unlike automatic variables, global variables retain their values between function calls, and are available to all functions. However, both types have initial values of zero.

Arrays or subscripted variables

An array, also referred to as a subscripted variable, is indicated with a single lowercase letter (the array name) followed by an expression in brackets (the subscript). For example,

`f[`*expression*`]`

The names of arrays can coincide with simple variable names or function names without conflicting. The subscript values must be greater than or equal to 0 and less than or equal to 2047; any fractional part of a subscript is discarded before use. Only one-dimensional arrays are permitted.

Subscripted variables can be used in expressions, function calls, and return statements. An array name can be used as an argument to a function or can be declared as automatic in a function definition by the use of empty brackets. For example:

```
f(a[])
define f(a[])
auto a[]
```

When an array name is declared automatic, the entire contents of the array is copied for the use of the function and thrown away on exit from the function. Such array names, used with empty brackets and referring to whole arrays, cannot be used in any context other than that just shown.

Statements

A statement is any direct instruction. Statements can be grouped together by surrounding them with braces, as in the body of a function definition:

```
define a(x) {
        statement
        statement ; statement
        return
    }
```

When statements are grouped, each individual statement must end with a semicolon or a newline to distinguish it from the next. Except where altered by control statements (such as a `while` loop), execution of grouped statements is sequential.

When a statement is an expression, the value of the expression is printed, followed by a newline character, unless the main operator is an assignment operator.

The following is a basic dictionary of `bc` predefined statements.

"*string*"
The quote statement prints the *string* contained within the quotation marks.

`break`
The `break` statement causes termination of a `for` or `while` statement.

`auto` *identifier*[, *identifier*] ...
The `auto` statement causes the values of one or more identifiers to be pushed down on the stack. The identifiers can be ordinary identifiers or array identifiers. Array identifiers are specified by following the array name with empty brackets. The `auto` statement must be the first statement in a function definition.

`define` *function-name*([*parameter*[, *parameter*] ...]) {*statements*}
The `define` statement defines a function. The parameters can be ordinary identifiers or array names. Array names must be followed by empty brackets.

`return`

`return(`*expression*`)`

The `return` statement causes the following actions:

- Termination of a function.
- Popping of the auto variables on the stack.
- Specifying the results of the function. The first form is equivalent to `return(0)`. The result of the function is the result of the expression in parentheses.

`quit`

The `quit` statement stops execution of a `bc` program and returns control to the A/UX system software when it is first encountered. Because it is not treated as an executable statement, it cannot be used in a function definition or in an `if`, `for`, or `while` statement.

`sqrt(`*expression*`)`

The result is the square root of the expression. The result is truncated in the least significant decimal place. The scale of the result is the scale of the expression or the value of `scale`, whichever is larger.

`length(`*expression*`)`

The result is the total number of significant decimal digits in the expression. The scale of the result is 0.

`scale(`*expression*`)`

The result is the number of the decimal point in the expression. The scale of the result is 0.

Assignment statements

`bc` assignment statements work in exactly the same manner as they do in the C programming language. Table 11-1 lists the assignment statement constructs.

Table 11-1 Assignment statements

`x=y=z`	Is the same as	`x=(y=z)`
`x =+y`	Is the same as	`x = x+y`
`x =-y`	Is the same as	`x = x-y`
`x = -y`	Is the same as	`x = -y`
`x =*y`	Is the same as	`x = x*y`
`x =/y`	Is the same as	`x = x/y`
`x =%y`	Is the same as	`x = x%y`
`x =^y`	Is the same as	`x = x^y`
`x++`	Is the same as	`(x=x+1)-1`
`x--`	Is the same as	`(x=x-1)+1`
`++x`	Is the same as	`x = x+1`
`--x`	Is the same as	`x = x-1`

◆ **Note** In some of these constructs, spaces are significant. There is an important difference between `x=-y` and `x= -y`. The first replaces `x` by `x-y` and the second replaces `x` by `-y`. ◆

All assignment operators are interpreted from right to left. The variables in an assignment statement should have single lowercase letter names. Ordinary variables are used as internal storage registers to hold integer values, and have an initial value of zero. The statement

```
x=x+3
```

has the effect of increasing by three the value of the contents of register `x`. In this case, although the increase in value is performed, that value is not printed. To print the value of `x` after the assignment, either explicitly call `x`, as in the following example:

```
x=x+3
x
```

or surround the assignment with parentheses. The latter instructs `bc` to treat the statement as the value of the result of the operation. The assignment can then be used anywhere an expression can be used. For example:

```
(x=x+3)
```

In this example, the value of `x` is incremented and the resulting value is printed.

The value of an assignment statement can be used even when it is not placed within parentheses. For example,

```
x=a[i=i+1]
```

instructs `bc` to increment `i` before using it as a subscript and then assign the resulting value to `x`.

Since each variable register name must be a unique, single lowercase letter, there can be only 26.

Control statements

The `if`, `while`, and `for` control statements are available in `bc` to alter the flow within programs or to cause iteration. They can be used individually as a simple statement or grouped to form a compound statement. A compound statement consists of a collection of statements enclosed in braces, as in a function definition.

Relational operators

Unlike all other operators, the `bc` relational operators are valid only as the object of an `if` or `while` statement or inside a `for` statement. Similarly, all control structures rely at least in part on the evaluation of a relational statement or expression. Table 11-2 illustrates the six relational operators and their definitions.

Table 11-2 Relational operators

Operator	Definition
<	Less than
>	Greater than
<=	Less than or equal to
>=	Greater than or equal to
==	Equal to
!=	Not equal to

◆ **Note** Do not use `=` instead of `==` as a relational operator. Both of these are legal, so there is no diagnostic message, but `=` does not do a comparison. The `=` operator is an assignment operator. ◆

The `if` statement

The `if` statement is a conditional statement that causes execution of its instruction if and only if the relation is true. Then, control passes to the next statement in sequence. The following is the standard format for an `if` statement in `bc`:

`if(`*relation*`)` *statement*

The `while` statement

`while` causes repeated execution of its instruction as long as the relation tests as true. The relation is tested before each execution of its range; if the result is true, the body of the `while` statement is executed, and the loop continues. If the relation is false, control passes to the next statement beyond the range of the `while` statement. The following format is standard for the `while` statement in `bc`:

```
while(relation) {
        statement
        statement
        ...
    }
```

The `for` statement

The typical use of a `for` statement is for controlled iteration. For example:

`for(`*expression1*`;` *relation*`;` *expression2*`)` *statements*

The `for` statement begins by executing *expression1*. Then the *relation* is tested. If the *relation* is true, the *statements* in the body of the `for` are executed. Then *expression2* is executed. The *relation* is then tested, and so forth, until the relational test fails.

The following example (in immediate mode) shows the proper use of the `for` statement. In this example, the function returns the factorial of the integer given as input:

```
define f(n) {
auto i, x
x=1
for(i=1; i<=n; i=i+1) x=x*i
return(x)
}
f(5)
120
f(3)
6
```

Expressions

The simplest `bc` expression is a single digit. An expression can consist of any number of operators and operands, provided that they represent a value.

The following points are important to remember when using expressions in `bc`:

- Any term in an expression can be preceded by a minus sign to indicate that it is a negative (the unary minus sign).
- The value of an expression is printed unless the main operator is an assignment.
- Division by zero produces an error comment.

Table 11-3 shows the operators that can be used in `bc` expressions, in order of precedence. Operators with the same precedence are grouped together.

Table 11-3 Operators and their precedence

Operator	Function
^	Exponentiation
*	Multiplication
%	Remaindering (integer result truncated toward 0)
/	Division
+	Addition
-	Subtraction
=	Assignment

Contents of parentheses are evaluated before items outside the parentheses. Exponentiations are performed from right to left, while the other operations are performed from left to right.

- `a^b^c` and `a^(b^c)` are equivalent
- `a-b*c` is the same as `a-(b*c)`
- `a/b*c` is equivalent to `(a/b)*c` because the expression is evaluated from left to right.

Brief descriptions of the various types of expressions recognized by `bc` are as follows :

-expression	The result is the negative of the *expression.*
++expression	The *expression* is incremented by one. The result is the value of the *expression* after incrementing.
--expression	The *expression* is decremented by one. The result is the value of the *expression* after decrementing.
expression ++	The *expression* is incremented by one. The result is the value of the *expression* before incrementing.
expression --	The *expression* is decremented by one. The result is the value of the *expression* before decrementing.
expression^*expression*	The result is the first *expression* raised to the power of the second *expression.* The second *expression* must be an integer. If *a* is the scale of the left expression and *b* is the absolute value of the right expression, the scale of the result is `min(`*a*`*`*b*`,max(scale,`*a*`))`

expression * *expression*

The result is the product of the two *expression* values. If *a* and *b* are the scales of the two expressions, the scale of the result is `min(`*a*`*`*b*`,max(scale,`*a*`,`*b*`))`

expression / *expression*

The result is the quotient of the two *expression* values. The scale of the result is the value of `scale`.

expression % *expression*

The `%` (modulus) operator produces the remainder of the division of the two *expression* values. More precisely, *a*%*b* has the same value as *a*-((*a*/*b*)*b*).

The scale of the result is the sum of the scales of the quotient and the divisor.

The additive operators bind left to right.

expression + *expression*

The result is the sum of the two *expression* values. The scale of the result is the maximum of the scales of the *expression* values.

expression - *expression*

The result is the difference of the two *expression* values. The scale of the result is the maximum of the scales of the *expression* values.

Input and output bases: `ibase` and `obase`

`bc` possesses a scaling provision that enables it to work in bases other than decimal. In addition, input and output can be set to different bases, for automatic conversion from one base to another. `ibase` handles the conversion for input, and `obase` for output.

`ibase` and `obase` have no effect on the course of internal computation or on the evaluation of expressions. They affect only input and output conversions, respectively.

`ibase`

The setting for `ibase` determines the base used for interpreting input, and is initially set to 10 (decimal). To set `ibase` to another base, use the `=` assignment operator. For example, the following assignment sets the input base to base 8:

```
ibase = 8
```

Assuming that the output base is set to decimal, with the `ibase` now set to octal, the input

```
11
```

automatically produces the following output:

```
9
```

If, at this point, you want to change the input base back to decimal, you must compensate for the fact that input is now being interpreted as octal. So, in setting the new base, you must use the correct octal value:

```
ibase = 12
```

Because the `ibase` is still set to octal, it interprets the 12 as an octal 10, and resets the base to decimal. Until reset again, `ibase` interprets all input in decimal.

For handling hexadecimal notation, the characters A through F are permitted in numbers (regardless of what base is in effect) and are interpreted as digits having values 10 through 15, respectively. The statement

```
ibase = A
```

changes the base to decimal regardless of the current input base.

`ibase` can handle base settings from 1 to 16. If larger or smaller settings are attempted, `ibase` disregards them. There is no error message to this effect and the last valid setting remains intact.

obase

The setting for `obase` is used for interpreting the output base and is initially set to 10 (decimal). Assuming that `ibase` is set to 10,

```
obase = 16
1000
```

produces the following output:

```
3E8
```

thus providing a simple decimal-to-hexadecimal conversion facility.

Very large output bases are permitted and are sometimes useful; for example, large numbers can be generated in groups of five digits by setting `obase` to 100000. Very large numbers are split across lines with 70 characters per line. To force the continuation of a line, end it with a backslash (\).

Decimal output conversion is practically instantaneous, but output of very large numbers (that is, more than 100 digits) with other bases is rather slow. Nondecimal output conversion of a 100-digit number takes about 3 seconds.

scale

The number of digits after the decimal point of a number is referred to as its scale. `bc` can handle numbers possessing up to 99 decimal places. The initial default setting for `scale` is 0. When the library option is invoked, however, the default is automatically set to 20. To set `scale` to a specific value, use the following statement:

`scale =` *n*

where *n* equals the new value of the `scale` setting. The contents of `scale` must be no greater than 99 and no less than its initial value of 0. However, appropriate scaling can be arranged when more than 99 fraction digits are required.

When two scaled numbers are combined by means of an arithmetic operation, the scale of the result is determined by the following rules:

Addition and subtraction
The scale of the result is the larger of the scales of the two operands. In this case, there is never any truncation of the result.

Multiplication
The scale of the result is never less than the maximum of the two scales of the operands and never more than the sum of the scales of the operands. Subject to those two restrictions, the scale of the result is set equal to the contents of the internal quantity `scale`.

Division
The scale of a quotient is the contents of the internal quantity `scale`. The scale of a remainder is the sum of the scales of the quotient and the divisor.

Exponentiation The result of an exponentiation is scaled as if the implied multiplications were performed. An exponent must be an integer.

Square root The scale of a square root is set to the maximum of the scale of the argument and the contents of `scale`.

All of the internal operations are actually carried out in terms of integers, with digits being discarded when necessary. In every case where digits are discarded, truncation (not rounding) is performed.

The value held in `scale` can be used in expressions just like other variables. The expression

```
scale = scale + 1
```

increases the value of `scale` by 1, and the statement

```
scale
```

causes the current value of `scale` to be printed.

It should be noted that, regardless of the `ibase` or `obase` settings, the `scale` setting is always interpreted in decimal base.

Part 4 Screen-Oriented Tools

A/UX also provides tools to control screen functions and to use dialog boxes for input and output.

You can use the `curses` package to write screen-oriented programs. `curses` provides a terminal-independent method of screen-oriented input and output. It includes facilities for taking input from the terminal, sending output to a terminal, creating and manipulating windows on the screen, and performing screen updates in an optimal fashion. A program using the `curses` routines and functions generally needs to know nothing about the capabilities of any particular terminal; these characteristics are determined at execution time and guide the program in taking input and producing output. Thus, programs using this package can interact with a large variety of terminals and terminal types.

The Commando tool is useful for screen-oriented input and output on Macintosh computers. Commando lets you create CommandShell command lines by selecting controls within Macintosh dialog boxes. Controls direct the placement of options on the command line. When the user selects a particular control, Commando places a specific option on the command line. Once they are constructed, the command lines are either placed in a CommandShell window for execution or executed in a subshell.

The following chapters detail these tools.

12 curses: Terminal-Independent Screen I/O

Overview of curses usage / 12-3
List of curses routines / 12-18
Operation details / 12-39
Example program: scatter / 12-47
Example program: show / 12-49
Example program: highlight / 12-51
Example program: window / 12-53
Example program: two / 12-55
Example program: termhl / 12-59
Example program: editor / 12-62

The curses package provides a terminal-independent method of providing screen-oriented input and output. It includes facilities for taking input from the terminal, sending output to a terminal, creating and manipulating windows on the screen, and performing screen updates in an optimal fashion. A program using the curses routines and functions generally needs to know nothing about the capabilities of any particular terminal; these characteristics are determined at execution time and guide the program in taking input and producing output. Thus, programs using this package can interact with a large variety of terminals and terminal types.

This chapter is an introduction to the `curses` and `terminfo` packages for writing screen-oriented programs. This chapter documents each `curses` function and discusses several sample programs. The sample programs are at the end of this chapter.

For older programs, `termcap` is provided for backward compatibility; new programs should use `terminfo`.

Overview of curses usage

For curses to be able to produce the proper output, it must know what kind of terminal you have. curses uses the standard A/UX system convention for this; the name of the terminal is stored in the environment variable TERM.

A program using curses always starts by calling initscr (see Listing 12-1). Other modes can then be set as needed by the program. Possible modes include cbreak and idlok(stdscr, TRUE). These modes are explained later.

A curses program follows the framework shown in Listing 12-1.

Listing 12-1 Framework of a curses program

```
#include <curses.h>
main()
{
...
initscr();        /* Initialization */
cbreak();         /* Various optional mode settings */
nonl();
noecho();
...
while (!done) { /* Main body of program */
   ...
    /* Sample calls to draw on screen */
    move(row, col);
    addch(ch);
    printw("Formatted print with value %d\n", value);
   ...
}
endwin();         /* Clean up */
exit(0);
}
```

Output

During the execution of the program, output to the screen is done with routines such as

`addch(`*ch*`)`

and

`printw(`*fmt, args*`)`

which behave just like `putchar` and `printf` except that they go through `curses`. The cursor can be moved with the call

`move(`*row, col*`)`

These routines generate output only to a data structure called a *window*, not to the actual screen. A window is a representation of a CRT screen, containing such things as an array of characters to be displayed on the screen, a cursor, a current set of video attributes, and various modes and options. Unless you use more than one of them, you don't need to worry about windows except to realize that a window is buffering your requests for output to the screen. For further information about windows, see the section "Multiple Windows" later in this chapter.

To send all accumulated output, you must call

```
refresh()
```

Finally, before the program terminates, it should call

```
endwin()
```

which restores all terminal settings and positions the cursor at the bottom of the screen.

See the sample program `scatter` at the end of this chapter. This program reads a file and displays it in a random order on the screen. Some programs assume all screens are 24 lines by 80 columns. It is important to understand that many are not. The variables

`LINES`

and

`COLS`

are defined by `initscr` with the current screen size. Programs should use them instead of assuming a 24 by 80 screen.

No output to the terminal actually happens until `refresh` is called. Instead, routines such as `move` and `addch` draw on a window data structure called `stdscr` (standard screen). `curses` always keeps track of what is on the physical screen, as well as what is in `stdscr`.

When `refresh` is called, `curses` compares the two screen images and sends a stream of characters to the terminal that turns the current screen into what is desired. `curses` considers many different ways to do this, taking into account the various capabilities of the terminal and similarities between what is on the screen and what is desired. It usually produces as few characters as is possible. This function is called *cursor optimization* and is the source of the name of the `curses` package.

◆ **Note** Because of the hardware scrolling of terminals, writing to the lower-right character position is impossible. ◆

Input

`curses` functions are also provided for input from the keyboard. The primary function is

```
getch()
```

which is like `getchar` except that it goes through `curses`. This function waits for the user to type a character on the keyboard and then returns that character. Its use is recommended for programs using the options

```
cbreak()
```

or

```
noecho()
```

because several terminal-dependent or system-dependent options become available that are not possible with `getchar`.

Options that you can use with `getch` include

```
keypad
```

which allows extra keys such as arrow keys, function keys, and other special keys that transmit escape sequences to be treated just as any other key. (The values returned for these keys are listed later; these values are over octal 400, so they should be stored in a variable larger than a `char`.)

The

```
nodelay
```

option causes the value -1 to be returned if there is no input waiting. Normally, `getch` waits until a character is typed.

Finally, the routine

`getstr(`*str*`)`

can be called, allowing input of an entire line, up to a newline. This routine handles echoing and the erase and kill characters of the user.

Examples of the use of these options are in later sample programs.

The following function keys might be returned by `getch`, if `keypad` is enabled. Note that not all of these can be supported by a particular terminal/keyboard, because the key doesn't exist or the terminal is not transmitting a unique code when the key is pressed.

Name	*Value*	*Key name*
`KEY_BREAK`	0401	Break key (unreliable)
`KEY_DOWN`	0402	The four arrow keys...
`KEY_UP`	0403	
`KEY_LEFT`	0404	
`KEY_RIGHT`	0405	
`KEY_HOME`	0406	Home key (upward + left arrow)
`KEY_BACKSPACE`	0407	Backspace (unreliable)
`KEY_F0`	0410	Function keys; space for 64 keys is reserved (only KF0 through KF10 are currently supported)
`KEY_F(n)`	`(KEY_F0+(n))`	Formula for `fn`
`KEY_DL`	0510	Delete line
`KEY_IL`	0511	Insert line
`KEY_DC`	0512	Delete character
`KEY_IC`	0513	Insert character or enter insert mode
`KEY_EIC`	0514	Exit from insert character mode
`KEY_CLEAR`	0515	Clear screen
`KEY_EOS`	0516	Clear to end of screen
`KEY_EOL`	0517	Clear to end of line

KEY_SF	0520	Scroll one line forward
KEY_SR	0521	Scroll one line backward (reverse)
KEY_NPAGE	0522	Next page
KEY_PPAGE	0523	Previous page
KEY_STAB	0524	Set tab
KEY_CTAB	0525	Clear tab
KEY_CATAB	0526	Clear all tabs
KEY_ENTER	0527	Enter or send (unreliable)
KEY_SRESET	0530	Soft (partial) reset (unreliable)
KEY_RESET	0531	Reset or hard reset (unreliable)
KEY_PRINT	0532	Print or copy
KEY_LL	0533	Home down or bottom (lower left)

The following keys are not currently supported on the Macintosh II: KEY_BREAK, KEY_ENTER, KEY_SRESET, KEY_RESET, and KEY_PRINT.

See the sample program show at the end of this chapter for an example of the use of getch. The show program pages through a file, showing one full screen each time the user presses the space bar. By creating an input file for show made up of 12-line pages, each segment varying slightly from the previous page, nearly any exercise for curses can be created. Such input files are called *show scripts.*

The following activities take place in the sample show program:

- cbreak is called so that you can press the space bar without having to press RETURN.
- noecho is called to prevent the space from echoing in the middle of a refresh, messing up the screen.
- nonl is called to enable more screen optimization.
- idlok is called to allow insert and delete lines, because many show scripts are constructed to duplicate bugs caused by that feature.
- clrtoeol clears from the cursor to the end of the line.
- clrtobot clears from the cursor to the end of the screen.

Highlighting

The function `addch` always draws two things on a window. In addition to the character itself, it draws a set of "attributes" associated with the character. These attributes cover various forms of highlighting of the character. For example, the character can be put in reverse video, bold, or underline. You can think of the attributes as the color of the ink used to draw the character.

A window always has a set of current attributes associated with it. The current attributes are associated with each character as it is written to the window. The current attributes can be changed with a call to

`attrset(`*attrs*`)`

(Think of this as dipping the window's pen in a particular color of ink.) The names of the attributes are

```
A_STANDOUT
A_REVERSE
A_BOLD
A_DIM
A_INVIS
A_UNDERLINE
```

For example, to put the word *boldface* in bold, you might use the following code:

```
printw("A word in ");
attrset(A_BOLD);
printw("boldface");
attrset(0);
printw(" really stands out.\n");
...
refresh();
```

Not all terminals are capable of displaying all attributes. If a particular terminal cannot display a requested attribute, `curses` attempts to find a substitute attribute. If none is possible, the attribute is ignored.

The A_STANDOUT attribute is used to make text attract the attention of the user. The particular hardware attribute used for A_STANDOUT varies from terminal to terminal, and is chosen to be the most visually pleasing attribute the terminal has. A_STANDOUT is typically implemented as reverse video or bold. Many programs don't really need a specific attribute, such as bold or inverse video, but instead just need to highlight some text. For such applications, the A_STANDOUT attribute is recommended. Two convenient functions,

```
standout()
standend()
```

turn this attribute on and off.

Attributes can be turned on in combination. For example, to turn on blinking bold text, use

```
attrset(A_BLINK|A_BOLD)
```

Individual attributes can be turned on and off with attron and ttroff without affecting other attributes.

For a sample program using attributes, see the highlight program at the end of this chapter. The highlight program takes a text file as input and allows embedded escape sequences to control attributes. In this sample program,

\U	turns on underlining
\B	turns on bold
\N	restores normal text

Note the initial call to scrollok. This allows the terminal to scroll if the file is longer than one screen. When an attempt is made to draw past the bottom of the screen, curses automatically scrolls the terminal up a line and calls refresh.

The highlight program comes as close to being a filter as is possible with curses. It is not a true filter, because curses must "take over" the CRT screen. To determine how to update the screen, it must know what is on the screen at all times. This requires curses to clear the screen in the first call to refresh and to know the cursor position and screen contents at all times.

Multiple windows

A window is a data structure representing all or part of the CRT screen. It has room for a two-dimensional array of characters, attributes for each character (a total of 16 bits per character: 7 for text and 9 for attributes), a cursor, a set of current attributes, and a number of flags.

`curses` provides a full screen window, called

`stdscr`

and a set of functions that use `stdscr`. Another window is provided called

`curscr`

representing the physical screen.

It is important to understand that a window is only a data structure. Use of more than one window does not imply use of more than one terminal, and it does not involve more than one process. A window is merely an object that can be copied to all or part of the terminal screen. The current implementation of `curses` does not allow windows that are bigger than the screen.

You can create additional windows with the function

`newwin(`*lines, cols, begin-row, begin-col*`)`

which returns a pointer to a newly created window. The window is *lines* by *cols*, and the upper-left corner of the window is at screen position (*begin-row*, *begin-col*).

All operations that affect `stdscr` have corresponding functions that affect an arbitrarily named window. Generally, these functions have names formed by putting a `w` on the front of the `stdscr` function and adding the window name as the first parameter. Thus,

`waddch(`*mywin*`, c)`

writes the character `c` to window *mywin*. The `wrefresh` function is used to flush the contents of a window to the screen.

Windows are useful for maintaining several different screen images, among which you can alternate. Also, you can subdivide the screen into several windows, refreshing each of them as desired. When windows overlap, the contents of the screen will be copied from the more recently refreshed window.

In all cases, the non-`w` version of the function calls the `w` version of the function, using `stdscr` as the additional argument. Thus, a call to

```
addch(c)
```

results in a call to

```
waddch(stdscr, c)
```

The sample program `window` at the end of this chapter shows the use of multiple windows. The main display is kept in `stdscr`. When the user temporarily wants to put something else on the screen, a new window is created covering part of the screen. A call to `wrefresh` on that window causes the window to be written over `stdscr` on the screen. Calling `refresh` on `stdscr` causes the original window to be redrawn on the screen.

In the sample `window` program, note the calls to

```
touchwin
```

before an overlapping window is written out. These are necessary to defeat an optimization in `curses`. If you have trouble refreshing a new window that overlaps an old window, it might be necessary to call `touchwin` on the new window to get it completely written out.

For convenience, a set of move functions are also provided for most of the common functions, which result in a call to `move` before the other function. For example:

`mvaddch(`*row*`,` *col*`, c)`

is the same as

`move(`*row*`,` *col*`); addch(c)`

Combinations also exist; for example,

`mvwaddch(`*row*`,` *col*`,` *win*`, c)`

Multiple terminals

`curses` can produce output on more than one terminal at once. This is useful for single-process programs that access a common database, such as multiplayer games. Output to multiple terminals is a difficult business, and `curses` does not solve all the problems for the programmer. The program itself must determine the filename of each terminal line and what kind of terminal is on each of those lines.

The standard method (checking `$TERM` in the environment) does not work because each process can examine only its own environment. Another problem that must be solved is that of multiple programs reading from one line. This situation produces a race condition and should be avoided. Nonetheless, a program wanting to take over another terminal cannot just shut off whatever program is currently running on that line. (Usually, security considerations also make this inappropriate. However, for some applications, such as an interterminal communication program or a program that takes over unused TTY lines, it is appropriate.)

A typical solution requires that the user logged in on each line run a program that notifies the master program that the user is interested in joining the master program, telling it the notification program process ID, the name of the TTY line, and the type of terminal being used. Then the program goes to sleep until the master program finishes. When done, the master program wakes up the notification program, and all programs exit.

`curses` handles multiple terminals by always having a "current terminal." All function calls always affect the current terminal. The master program should set up each terminal, saving a reference to the terminals in its own variables. When it wishes to affect a terminal, it should set the current terminal as desired, and then call ordinary `curses` routines.

References to terminals have type `struct screen *`.

A new terminal is initialized by calling

`newterm(`*type*`,`*fd*`)`

`newterm` returns a screen reference to the terminal being set up; *type* is a character string, naming the kind of terminal being used; and *fd* is a `stdio` file descriptor to be used for input and output to the terminal. (If only output is needed, the file can be open for output only.)

This call replaces the normal call to `initscr`, which calls

```
newterm(getenv("TERM"),stdout)
```

To change the current terminal, call

`set_term(`*sp*`)`

where *sp* is the screen reference to be made current. `set_term` returns a reference to the previous terminal.

It is important to realize that each terminal has its own set of windows and options. Each terminal must be initialized separately with `newterm`. Options such as `cbreak` and `noecho` must be set separately for each terminal. The functions `endwin` and `refresh` must be called separately for each terminal. See Figure 12-2 for a typical scenario to send a message to each terminal.

Listing 12-2 Sending a message to several terminals

```
for (i=0; i<nterm; i++) {
    set_term(terms[i]);
    mvaddstr(0, 0, "Important message");
    refresh();
}
```

See the sample program `two` at the end of this chapter for a full illustration. The `two` program pages through a file, showing one page to the first terminal and the next page to the second terminal. It then waits for a space to be typed on either terminal, and shows the next page to the terminal typing the space. Each terminal must be separately put into `nodelay` mode. It is necessary to busy-wait or call `sleep` (see `sleep`(3C) in *A/UX Programmer's Reference*) between each check for keyboard input, or use the multiplexor `select`(2). This program sleeps for a second between checks.

The `two` program is just a simple example of two-terminal `curses`. It does not handle notification, as described earlier; instead, it requires the name and type of the second terminal on the command line. As written, the command

```
sleep 100000
```

must be typed on the second terminal to put it to sleep while the program runs, and the first user must have both read and write permission on the second terminal.

Low-level `terminfo` usage

Some programs need to use lower-level primitives than those offered by `curses`. For such programs, the `terminfo`-level interface is offered.

The `terminfo`-level interface does not manage your CRT screen, but rather gives you access to strings and capabilities that you can use to manipulate the terminal. `curses` takes care of all the glitches and odd features present in physical terminals, but at the `terminfo` level you must deal with them yourself. Whenever possible, the higher-level `curses` routines should be used. This makes your program more portable to other A/UX systems and to a wider class of terminals. Also, it cannot be guaranteed that this part of the interface does not change or is upwardly compatible with previous releases.

There are two circumstances in which you should use `terminfo`. The first is when you are writing a special-purpose tool that sends a special-purpose string to the terminal, such as programming a function key, setting tab stops, sending output to a printer port, or dealing with the status line. The second situation is when you are writing a filter. A typical filter does one transformation on the input stream without clearing the screen or addressing the cursor. If this transformation is terminal-dependent and clearing the screen is inappropriate, use `terminfo`.

A program written at the `terminfo` level uses the framework shown in Figure 12-3.

Initialization is done by calling `setupterm`.

Passing the values 0, 1, 0 invokes reasonable defaults. If `setupterm` cannot figure out what kind of terminal you are on, it prints an error message and quits. Your program should call `reset_shell_mode` before it exits. Global variables with names like `clear_screen` and `cursor_address` are initialized by the call to `setupterm`. They can be produced using `putp` or `tputs` (which allows the programmer more control). These strings should not be directly sent to the terminal using `printf`, because they contain padding information. A program that directly generates strings fails on terminals that require padding or that use the xon/xoff flow-control protocol.

In the `terminfo` level, the higher-level routines described previously are not available. It is up to the programmer to generate whatever is needed. For a list of capabilities and a description of what they do, see `terminfo`(4).

Listing 12-3 `terminfo`-level framework

```
#include <curses.h>
#include <term.h>
    ...
    setupterm(0, 1, 0);
    ...
    putp(clear_screen);
    ...
    reset_shell_mode();
exit(0);
```

The `termhl` sample program at the end of this chapter shows a simple use of `terminfo`. It is a version of the `highlight` sample program that uses `terminfo` instead of `curses`. This version can be used as a filter. The strings to enter bold and underline mode, and to turn off all attributes, are used.

This program is more complex than it has to be to illustrate some properties of `terminfo`. The routine `vidattr` could have been used instead of directly generating

```
enter_bold_mode
enter_underline_mode
exit_attribute_mode
```

In fact, the program would be more robust if it did so, since there are several ways to change video attribute modes. However, this program was written to illustrate typical use of `terminfo`.

The function

`tputs(`*cap, affcnt, outc*`)`

applies padding information. Some capabilities contain strings like `$<20>`. This means to pad for 20 milliseconds. `tputs` generates enough pad characters to delay for the appropriate time. The first parameter is the string capability to be generated. The second is the number of lines affected by the capability. Some capabilities might require padding that depends on the number of lines affected. For example, `insert_line` might have to copy all lines below the current line, and might require time proportional to the number of lines copied. By convention, *affcnt* is 1 if no lines are affected. For safety, the value 1 is used rather than 0 (*affcnt* is multiplied by the amount of time per item, and anything multiplied by 0 is 0). The third parameter is a routine to be called with each character.

For many simple programs, *affcnt* is always 1 and *outc* always just calls `putchar`. For these programs, the routine `putp(`*cap*`)` is a convenient abbreviation. The `termhl` sample program can be simplified by using `putp`.

Note also in this example the special check for the capability `underline_char`. Some terminals, rather than having a code to start underlining and a code to stop underlining, have a code to underline the current character. The `termhl` program keeps track of the current mode, and if the current character is supposed to be underlined, outputs `underline_char` if necessary.

Low-level details such as this are precisely why the `curses` level is recommended over the `terminfo` level. `curses` takes care of terminals with different methods of underlining and other CRT functions. Programs at the `terminfo` level must handle such details themselves.

A larger example

For a final example, see the `editor` sample program at the end of this chapter.

The `editor` program illustrates how to use `curses` to write a screen editor patterned after the `vi` editor. This editor keeps the buffer in `stdscr` to keep the program simple; a real screen editor keeps a separate data structure. Many simplifications have been made here. No provision is made for files of any length other than the size of the screen, for lines longer than the width of the screen, or for control characters in the file.

Several points about this program are worth noting. The routine to write out the file illustrates the use of the `mvinch` function, which returns the character in a window at a given position. The data structure used here does not have a provision for keeping track of the number of characters in a line, or the number of lines in the file, so trailing blanks are eliminated when the file is written out.

The program uses these built-in `curses` functions:

```
insch
delch
insertln
deleteln
```

These functions behave much as the similar functions on intelligent terminals behave, inserting and deleting a character or a line.

The command interpreter accepts not only ASCII characters, but also special keys. (Some editors are "modeless," using nonprinting characters for commands. This is largely a matter of taste; the point being made here is that both arrow keys and ordinary ASCII characters should be handled.)

In the `editor` sample program, note the call to `mvaddstr` in the input routine. `addstr` is roughly like the C `fputs` function, which writes out a string of characters. Like `fputs`, `addstr` does not add a trailing newline. It is the same as a series of calls to `addch` using the characters in the string. `mvaddstr` is the `mv` version of `addstr`, which moves to the given location in the window before writing.

The CONTROL-L command illustrates a feature that most programs using `curses` should add. Often some program beyond the control of `curses` has written something to the screen, or some line noise has messed up the screen beyond what `curses` can keep track of. In this case, the user types CONTROL-L, causing the screen to be cleared and redrawn. This is done with the call to

```
clearok(curscr)
```

which sets a flag causing the next `refresh` to first clear the screen. Then `refresh` is called to force the redraw.

Note also the call to

```
flash()
```

which flashes the screen, if possible, and otherwise rings the bell. Flashing the screen is intended as a bell replacement, and is particularly useful if the bell bothers someone within earshot of the user. The routine

```
beep()
```

can be called when a real beep is desired. (If, for some reason, the terminal is unable to beep but able to flash, a call to `beep` flashes the screen.)

Another important point is that the input command is terminated by CONTROL-D, not ESCAPE. It is very tempting to use ESCAPE as a command, because ESCAPE is one of the few special keys that is available on most keyboards. (RETURN and BREAK are among the others.) However, using ESCAPE as a separate key introduces an ambiguity. Most terminals use sequences of characters beginning with ESCAPE (escape sequences) to control the terminal, and have special keys that send escape sequences to the computer. If the computer recognizes an ESCAPE coming from the terminal, it cannot determine whether the user pressed the ESCAPE key, or whether a special key was pressed. `curses` handles the ambiguity by waiting for up to 1 second. If another character is received during this second, and if that character might be the beginning of a special key, more input is read (waiting for up to 1 second for each character) until either (1) a full special key is read, (2) 1 second passes, or (3) a character is received that cannot have been generated by a special key.

While this strategy works most of the time, it is not foolproof. It is possible for the user to press ESCAPE, then to type another key quickly, which causes `curses` to think a special key has been pressed. Also, there is a 1-second pause until the escape can be passed to the user program, resulting in slower response to the ESCAPE key.

Many existing programs use ESCAPE as a fundamental command, so it cannot be changed without infuriating a large class of users. Such programs cannot make use of special keys without dealing with this ambiguity, and at best must resort to a timeout solution. The message is clear: When designing your program, avoid the ESCAPE key.

List of `curses` routines

This section describes all the routines available to the programmer in the `curses` package. The routines are organized by function. For an alphabetical list, see `curses`(3X).

Structure

All programs using `curses` should include the file `<curses.h>`. This file defines several `curses` functions as macros, and defines several global variables and the datatype `WINDOW`. References to windows are always of type `WINDOW *`.

`curses` also defines certain windows as constants:

`stdscr`	the standard screen, used as a default to routines expecting a window
`curscr`	the current screen, used only for certain low-level operations like clearing and redrawing a garbaged screen

Integer variables are declared, containing the size of the screen:

`LINES`	number of lines on the screen
`COLS`	number of columns on the screen

Boolean constants are defined as follows with values 1 and 0:

```
#define TRUE(1)
#define FALSE(0)
#define ERR(-1)
#define OK(0)
```

Additional constants are values returned from most `curses` functions:

`ERR`	returned if there was some error, such as moving the cursor outside a window
`OK`	returned if the function was properly completed

The include file `<curses.h>` automatically includes `<stdio.h>` and an appropriate TTY driver interface file, currently either `<sgtty.h>` or `<termio.h>`.

◆ **Note** The driver interface `<sgtty.h>` is a TTY driver interface used in other versions of the UNIX system. ◆

Including `<stdio.h>` again is harmless but wasteful; including `<sgtty.h>` again usually results in a fatal error.

A program using `curses` should include the loader option

```
-lcurses
```

in the makefile. This is true both for the `terminfo` level and the `curses` level.

The compilation flag

```
-DMINICURSES
```

can be included if you restrict your program to a small subset of `curses` concerned primarily with screen output and optimization. The routines possible with `mini-curses` are listed in the section "`mini-curses`" later in this chapter.

Initialization

The following functions are called when initializing a program:

`initscr()`
The first function called should always be `initscr`. This determines the terminal type and initializes `curses` data structures. `initscr` also arranges that the first call to `refresh` clears the screen.

`endwin()`
A program should always call `endwin` before terminating. This function restores TTY modes, moves the cursor to the lower-left corner of the screen, resets the terminal into the proper nonvisual mode, and tears down all appropriate data structures.

`newterm(`*type*, *fd*`)`
A program that generates output to more than one terminal should use `newterm` instead of `initscr`. `newterm` should be called once for each terminal. It returns a variable of type `SCREEN *` which should be saved as a reference to that terminal. The arguments are the type of the terminal (a string) and a `stdio` file descriptor `(FILE *)` for output to the terminal. The file descriptor should be open for both reading and writing if input from the terminal is desired. The program should also call `endwin` for each terminal being used (see `set_term`). If an error occurs, the value NULL is returned.

`set_term(`*new*`)`
This function is used to switch to a different terminal. The screen reference *new* becomes the new current terminal. The previous terminal is returned by the function. All other calls affect only the current terminal.

`longname()`
This function returns a pointer to a static area containing a verbose description of the current terminal. It is defined only after a call to `initscr`, `newterm`, or `setupterm`.

Option setting

The functions described here set options within `curses`. In each case, *win* is the window affected, and *bf* is a Boolean flag with value `TRUE` or `FALSE` (indicating whether to enable or disable the option). All options are initially `FALSE`. It is not necessary to turn these options off before calling `endwin`.

`clearok(`*win*, *bf*`)`
If set, the next call to `wrefresh` with this window clears the screen and redraws the entire screen. If *win* is `curscr`, the next call to `wrefresh` with any window causes the screen to be cleared. This is useful when the contents of the screen are uncertain, or in some cases for a more pleasing visual effect.

`idlok(`*win, bf*`)`
If this feature is enabled, `curses` considers using the hardware insert/delete line feature of terminals so equipped. If disabled, `curses` never uses this feature. The insert/delete character feature is always considered. Enable this option only if your application needs insert/delete line, for example, for a screen editor. It is disabled by default because insert/delete line tends to be visually annoying when used in applications where it is not really needed. If insert/delete line cannot be used, `curses` redraws the changed portions of all lines that do not match the desired line.

`keypad(`*win, bf*`)`
This option enables the keypad of the user's terminal. If enabled, the user can press a function key (such as an arrow key) and `getch` returns a single value representing the function key. If `keypad` is disabled, `curses` does not treat function keys specially. If the keypad in the terminal can be turned on (made to transmit) and off (made to work locally), turning on this option turns on the terminal keypad.

`leaveok(`*win, bf*`)`
Normally, the hardware cursor is left at the location of the window cursor being refreshed. This option allows the cursor to be left wherever the update happens to leave it. It is useful for applications where the cursor is not used because it reduces the need for cursor motions. If possible, the cursor is made invisible when this option is enabled.

`meta(`*win, bf*`)`
If enabled, characters returned by `getch` are transmitted with all 8 bits, instead of stripping the highest bit. The value `OK` is returned if the request succeeded; the value `ERR` is returned if the terminal or system is not capable of 8-bit input.

`meta` mode is useful for extending the nontext command set in applications where the terminal has a meta shift key. `curses` takes whatever measures are necessary to arrange for 8-bit input. On other versions of UNIX systems, `raw` mode is used. On A/UX systems, the character size is set to 8, parity checking disabled, and stripping of the eighth bit turned off.

Note that 8-bit input is a fragile mode. Many programs and networks pass only 7 bits. If any link in the chain from the terminal to the application program strips the eighth bit, 8-bit input is impossible.

`nodelay(win, bf)`
This option causes `getch` to be a nonblocking call. If no input is ready, `getch` returns -1. If disabled, `getch` hangs until a key is pressed.

`intrflush(win, bf)`
If this option is enabled when an interrupt key is pressed on the keyboard (interrupt, quit, suspend), all output in the TTY driver queue is flushed, giving the effect of faster response to the interrupt but causing `curses` to have the wrong idea of what is on the screen. Disabling the option prevents the flush. The default is for the option to be enabled. This option depends on support in the underlying teletype driver.

`typeahead(fd)`
Sets the file descriptor for typeahead check. *fd* should be an integer returned from `open` or `fileno`. Setting `typeahead` to -1 disables typeahead check. By default, file descriptor 0 (`stdin`) is used. `typeahead` is checked independently for each screen, and for multiple interactive terminals it should probably be set to the appropriate input for each screen. A call to `typeahead` always affects only the current screen.

`scrollok(win, bf)`
This option controls what happens when the cursor of a window is moved off the edge of the window, either from a newline on the bottom line or because the last character of the last line was typed. If disabled, the cursor is left on the bottom line. If enabled, `wrefresh` is called on the window, and then the physical terminal and window are scrolled up one line. Note that to get the physical scrolling effect on the terminal, it is also necessary to call `idlok`.

`setscrreg(t, b)`
`wsetscrreg(wing, t, bf)`
These two functions allow the user to set a software scrolling region in a window *win* or `stdscr`. *t* and *b* are the line numbers of the top and bottom margins of the scrolling region. (Line 0 is the top line of the window.) If this option and `scrollok` are enabled, an attempt to move off the bottom margin line causes all lines in the scrolling region to scroll up one line. Note that this has nothing to do with use of a physical scrolling region capability in the terminal, like that in the VT100. Only the text of the window is scrolled. If `idlok` is enabled and the terminal has either a scrolling region or insert/delete-line capability, they are probably used by the output routines.

Terminal mode setting

The functions described here are used to set modes in the TTY driver. The initial mode usually depends on the setting when the program is called; the initial modes documented here represent the normal situation.

`cbreak()`
`nocbreak()`
These two functions put the terminal into and out of `CBREAK` mode. In this mode, characters typed by the user are immediately available to the program. When out of this mode, the teletype driver buffers characters typed until newline is typed. Interrupt and flow-control characters are unaffected by this mode. Initially the terminal is not in `CBREAK` mode. Most interactive programs using `curses` set this mode.

`echo()`
`noecho()`
These functions control whether characters typed by the user are echoed as typed. Initially, characters typed are echoed by the teletype driver. Authors of many interactive programs prefer to do their own echoing in a controlled area of the screen, or not to echo at all, so they disable echoing.

`nl()`
`nonl()`
These functions control whether newline is translated into carriage return and line feed on output, and whether return is translated into newline on input. Initially, the translations do occur. By disabling these translations, `curses` is able to make better use of the line-feed capability, resulting in faster cursor motion.

`raw()`
`noraw()`
The terminal is placed into or out of `raw` mode. `raw` mode is similar to `cbreak` mode in that characters typed are immediately passed through to the user program. The differences are that in `raw` mode, the interrupt, quit, and suspend characters are passed through uninterpreted instead of generating a signal. `raw` mode also causes 8-bit input and output. The behavior of the BREAK key might be different on different systems.

`resetty()`
`savetty()`
These functions save and restore the state of the TTY modes. `savetty` saves the current state in a buffer; `resetty` restores the state to what it was at the last call to `savetty`.

Window manipulation

`newwin(`*num-lines*, *num-cols*, *beg-row*, *beg-col*`)`
Creates a new window with the given number of lines and columns. The upper-left corner of the window is at line *beg-row* column *beg-col*. If either *num-lines* or *num-cols* is 0, they default to `LINES-`*beg-row* and `COLS-`*beg-col*. A new full-screen window is created by calling `newwin(0,0,0,0)`.

`newpad(`*num-lines*, *num-cols*`)`
Creates a new `pad` data structure. A pad is like a window, except that it is not restricted by the screen size and is not associated with a particular part of the screen. Pads can be used when a large window is needed, and only a part of the window is on the screen at one time. Automatic refreshes of pads (for example, from scrolling or echoing of input) do not occur. It is not legal to call `refresh` with a pad as an argument; the routines `prefresh` or `pnoutrefresh` should be called instead. Note that these routines require additional parameters to specify the part of the pad to be displayed and the location on the screen to be used for display.

`subwin(`*orig*, *num-lines*, *num-cols*, *begy*, *begx*`)`
Creates a new window with the given number of lines and columns. The window is at position (*begy, begx*) on the screen. (It is relative to the screen, not *orig*.) The window is made in the middle of the window *orig*, so that changes made to one window affect both windows. When using this function, it is often necessary to call `touchwin` before calling `wrefresh`.

`delwin(win)`
Deletes the named window, freeing all memory associated with it. In the case of overlapping windows, subwindows should be deleted before the main window.

`mvwin(win, br, bc)`
Moves the window so that the upper-left corner is at position (*br*, *bc*). If the move would cause the window to be off the screen, it is an error and the window is not moved.

`touchwin(win)`
Throws away all optimization information about which parts of the window have been touched, by pretending the entire window has been drawn on. This is sometimes necessary when using overlapping windows, because a change to one window affects the other window, but the records of which lines have been changed in the other window does not reflect the change.

`overlay(win1, win2)`
`overwrite(win1, win2)`
These functions overlay *win1* on top of *win2*; that is, all text in *win1* is copied into *win2*. The difference is that `overlay` is nondestructive (blanks are not copied) and `overwrite` is destructive.

Causing output to the terminal

`refresh()`
`wrefresh(win)`
These functions must be called to get any output on the terminal, as other routines merely manipulate data structures. `wrefresh` copies the named window to the physical terminal screen, taking into account what is already there to do optimizations. `refresh` is the same, using `stdscr` as a default screen. Unless `leaveok` is enabled, the physical cursor of the terminal is left at the location of the window cursor.

`doupdate()`

`wnoutrefresh(`*win*`)`

These two functions allow multiple updates with more efficiency than `wrefresh`. To use them, it is important to understand how `curses` works. In addition to all the window structures, `curses` keeps two data structures representing the terminal screen: a "physical" screen, describing what is actually on the screen, and a "virtual" screen, describing what the programmer *wants* to have on the screen. `wrefresh` works by first copying the named window to the virtual screen (`wnoutrefresh`), and then calling the routine to update the screen (`doupdate`). If the programmer wishes to produce several windows at once, a series of calls to `wrefresh` results in alternating calls to `wnoutrefresh` and `doupdate`, causing several bursts of output to the screen. By calling `wnoutrefresh` for each window, it is then possible to call `doupdate` once, resulting in only one burst of output, with probably fewer total characters transmitted.

`prefresh(`*pad, pminrow, pmincol*
sminrow, smincol
smaxrow, smaxcol`)`

`pnoutrefresh(`*pad, pminrow, pmincol*
sminrow, smincol
smaxrow, smaxcol`)`

These routines are analogous to `wrefresh` and `wnoutrefresh` except that pads, instead of windows, are involved. The additional parameters are needed to indicate what part of the pad and screen are involved. *pminrow* and *pmincol* specify the upper-left corner, in the pad, of the rectangle to be displayed. *sminrow*, *smincol*, *smaxrow*, and *smaxcol* specify the edges, on the screen, of the rectangle to be displayed. The lower-right corner in the pad of the rectangle to be displayed is calculated from the screen coordinates, because the rectangles must be the same size. Both rectangles must be entirely contained within their respective structures.

Writing on window structures

The routines described here are used to "draw" text on windows. In all cases, a missing *win* is taken to be `stdscr`. *y* and *x* are the row and column, respectively. The upper-left corner is always (0,0), not (1,1). The `mv` functions imply a call to `move` before the call to the other function.

Moving the cursor

`move(`*y*`,`*x*`)`
`wmove(`*win*`,`*y*`,`*x*`)`
The cursor associated with the window is moved to the given location. This does not move the physical cursor of the terminal until `refresh` is called. The position specified is relative to the upper-left corner of the window. The position specified is relative to the screen, not to the individual window. Thus, if you have a window that is not in the upper-left corner of the screen, moving to the upper-left corner of the window would require the screen coordinates of that corner of the window rather than (0,0) to be passed to move.

Writing one character

`addch(`*ch*`)`
`waddch(`*win*`,`*ch*`)`
`mvaddch(`*y*`,`*x*`,`*ch*`)`
`mvwaddch(`*win*`,`*y*`,`*x*`,`*ch*`)`
The character *ch* is put in the window at the current cursor position of the window. If *ch* is a tab, newline, or backspace, the cursor is moved appropriately in the window. If *ch* is a different control character, it is drawn in the ^X notation. The position of the window cursor is advanced. At the right margin, an automatic newline is performed. At the bottom of the scrolling region, if `scrollok` is enabled, the scrolling region is scrolled up one line.

The *ch* parameter is actually an integer, not a character. Video attributes can be combined with a character by ORing them into the parameter. This results in these attributes also being set. (The intent here is that text, including attributes, can be copied from one place to another with `inch` and `addch`.)

Writing a string

`addstr(`*str*`)`
`waddstr(`*win*`,` *str*`)`
`mvaddstr(`*y*`,` *x*`,` *str*`)`
`mvwaddstr(`*win*`,` *y*`,` *x*`,` *str*`)`
These functions write all the characters of the null terminated character string *str* on the given window. They are identical to a series of calls to `addch`.

Clearing areas of the screen

`erase()`
`werase(`*win*`)`
These functions copy blanks to every position in the window.

`clear()`
`wclear(`*win*`)`
These functions are like `erase` and `werase` but they also call `clearok`, arranging that the screen is cleared on the next call to `refresh` for that window.

`clrtobot()`
`wclrtobot(`*win*`)`
All lines below the cursor in this window are erased. Also, the current line to the right of the cursor is erased.

`clrtoeol()`
`wclrtoeol(`*win*`)`
The current line to the right of the cursor is erased.

Inserting and deleting text

`delch()`
`wdelch(`*win*`)`
`mvdelch(`*y*`,`*x*`)`
`mvwdelch(`*win*`,`*y*`,`*x*`)`
The character under the cursor in the window is deleted. All characters to the right on the same line are moved to the left one position. This does not imply use of the hardware delete-character feature.

`deleteln()`
`wdeleteln(`*win*`)`
The line under the cursor in the window is deleted. All lines below the current line are moved up one line. The bottom line of the window is cleared. This does not imply use of the hardware delete-line feature.

`insch(`*c*`)`
`winsch(`*win*`,`*c*`)`
`mvinsch(`*y*`,`*x*`,`*c*`)`
`mvwinsch(`*win*`,`*y*`,`*x*`,`*c*`)`
The character *c* is inserted before the character under the cursor. All characters to the right are moved one space to the right, possibly losing the right-most character on the line. This does not imply use of the hardware insert-character feature.

`insertln()`
`winsertln(`*win*`)`
A blank line is inserted above the current line. The bottom line is lost. This does not imply use of the hardware insert-line feature.

Formatted output

`printw(`*fmt, args*`)`
`wprintw(`*win, fmt, args*`)`
`mvprintw(`*y, x, fmt, args*`)`
`mvwprintw(`*win, y, x, fmt, args*`)`
These functions correspond to `printf`. The characters that would be produced by `printf` are instead produced using `waddch` on the given window.

Miscellaneous

`box(`*win, vert, hor*`)`
A box is drawn around the edge of the window. *vert* and *hor* are the characters with which the box is to be drawn.

`scroll(`*win*`)`
The window is scrolled up one line. This involves moving the lines in the window data structure. As an optimization, if the window is `stdscr` and the scrolling region is the entire window, the physical screen is scrolled at the same time.

Input from a window

`getyx(`*win, y, x*`)`
The cursor position of the window is placed in the two integer variables *y* and *x*. Since this is a macro, no `&` is necessary for *x* or *y*.

`inch()`
`winch(`*win*`)`
`mvinch(`*y, x*`)`
`mvwinch(`*win, y, x*`)`
The character at the current position in the named window is returned. If any attributes are set for that position, their values are ORed into the value returned. The predefined constants `A_ATTRIBUTES` and `A_CHARTEXT` can be used with the `&` operator to extract the character or attributes alone. For example:

```
#include <curses.h>
...
        char c;
        ...
        c = inch() & A_CHARTEXT;
        ...
```

Input from the terminal

`getch()`
`wgetch(`*win*`)`
`mvgetch(`*y*`,`*x*`)`
`mvwgetch(`*win*`,`*y*`,`*x*`)`
A character is read from the terminal associated with the window. In `nodelay` mode, if there is no input waiting, the value -1 is returned. In `delay` mode, the program hangs until the system passes text through to the program. Depending on the setting of `cbreak`, this is after one character, or after the first newline.

If `keypad` mode is enabled, and a function key is pressed, the code for that function key is returned instead of the raw characters. Possible function keys are defined with integers beginning with 0401, whose names begin with `KEY_`. These are listed in the section "Input" earlier in this chapter. If a character is received that might be the beginning of a function key (such as ESCAPE), `curses` sets a 1-second timer. If the remainder of the sequence does not come within 1 second, the character is passed through; otherwise the function key value is returned. For this reason, on many terminals there is a 1-second delay after a user presses the ESCAPE key. (Using the ESCAPE key for a single character function is discouraged.)

`getstr(`*str*`)`
`wgetstr(`*win*`,`*str*`)`
`mvgetstr(`*y*`,`*x*`,`*str*`)`
`mvwgetstr(`*win*`,`*y*`,`*x*`,`*str*`)`
A series of calls to `getch` is made, until a newline is received. The resulting value is placed in the area pointed at by the character pointer *str*. The user's erase and kill characters are interpreted.

`scanw(`*fmt, args*`)`
`wscanw(`*win, fmt, args*`)`
`mvscanw(`*y, x, fmt, args*`)`
`mvwscanw(`*win, y, x, fmt, args*`)`
This function corresponds to `scanf`. `wgetstr` is called on the window, and the resulting line is used as input for the scan.

Video attributes

`attroff(`*at*`)`
`wattroff(`*win, attrs*`)`
`attron(`*at*`)`
`wattron(`*win, attrs*`)`
`attrset(`*at*`)`
`wattrset(`*win, attrs*`)`
`standout()`
`standend()`
`wstandout(`*win*`)`
`wstandend(`*win*`)`
These functions set the current attributes of the named window. These attributes can be any combination of `A_STANDOUT`, `A_REVERSE`, `A_BOLD`, `A_DIM`, `A_BLINK`, and `A_UNDERLINE`. These constants are defined in `<curses.h>` and can be combined with the C language OR operator (`|`).

The current attributes of a window are applied to all characters that are written into the window with `waddch`. Attributes are a property of the character and move with the character through any scrolling and insert/delete line/character operations. To the extent possible on the particular terminal, they are displayed as the graphic rendition of characters put on the screen.

`attrset(`*at*`)`

sets the current attributes of the given window to *at*.

`attroff(`*at*`)`

turns off the named attributes without affecting any other attributes.

`attron(at)`

turns on the named attributes without affecting any others.

`standout`

is the same as `attron(A_STANDOUT)`.

`standend`

is the same as `attrset(0)`; that is, it turns off all attributes.

Bells and flashing lights

`beep()`

`flash()`

These functions are used to signal the user. `beep` sounds the audible alarm on the terminal, if possible: if not, it flashes the screen (visible bell), if that is possible. `flash` flashes the screen or, if that is not possible, sounds the audible signal. If neither signal is possible, nothing happens. Nearly all terminals have an audible signal (a bell or beep) but only some can flash the screen.

Portability functions

The functions described here do not directly involve terminal-dependent character output but tend to be needed by programs that use `curses`. Unfortunately, their implementation varies from one version of UNIX to another. They are included here to enhance the portability of programs using `curses`.

`baudrate()`

`baudrate` returns the output speed of the terminal. The number returned is the integer baud rate—for example, 9600, rather than a table index such as `B9600`.

`erasechar()`

The erase character chosen by the user is returned. This is the character typed by the user to erase the character just typed.

`killchar()`
The line-kill character chosen by the user is returned. This is the character typed by the user to abort the entire line being typed.

`flushinp()`
This function throws away any typeahead that was typed by the user but not yet read by the program.

Delays

The functions described here are highly unportable, but are often needed by programs that use `curses`, especially real-time response programs. Some of these functions require a particular operating system or a modification to the operating system to work. In all cases, the routine compiles and returns an error status if the requested action is not possible. It is recommended that you avoid use of these functions if possible.

`draino(`*ms*`)`
The program is suspended until the output queue has drained enough to complete in *ms* additional milliseconds. Thus,

```
draino(50)
```

at 1200 baud pauses until there are no more than six characters in the output queue, because it takes 50 milliseconds to output the additional six characters. The purpose of this routine is to keep the program (and thus the keyboard) from getting ahead of the screen. If the operating system does not support the `ioctls` needed to implement `draino`, the value `ERR` is returned; otherwise, `OK` is returned.

`napms(`*ms*`)`
This function suspends the program for *ms* milliseconds. It is similar to `sleep` except with higher resolution. The resolution actually provided varies with the facilities available in the operating system, and often a change to the operating system is necessary to produce good results. If resolution of at least .1 second is not possible, the routine rounds to the next higher second, calls `sleep`, and returns `ERR`. Otherwise, the value `OK` is returned. Often the resolution provided is 1/60th second.

Lower-level functions

The functions described here are provided for programs not needing the screen optimization capabilities of `curses`. Programs are discouraged from working at this level, because they must handle various glitches in certain terminals. However, a program can be smaller if it only brings in the low-level routines.

Cursor motion

`mvcur(`*oldrow, oldcol, newrow, newcol*`)`
This routine optimally moves the cursor from (*oldrow, oldcol*) to (*newrow, newcol*). The user program is expected to keep track of the current cursor position. Note that unless a full screen image is kept, `curses` must make pessimistic assumptions, sometimes resulting in less than optimal cursor motion. For example, moving the cursor a few spaces to the right can be done by transmitting the characters being moved over, but if `curses` does not have access to the screen image, it does not know what these characters are.

`terminfo` level

These routines are called by low-level programs that need access to specific capabilities of `terminfo`. A program working at this level should include both `<curses.h>` and `<term.h>`, in that order. After a call to `setupterm`, the capabilities are available with macro names defined in `<term.h>`. See `terminfo`(4) for a detailed description of the capabilities.

Boolean-valued capabilities have the value 1 if the capability is present and 0 if it is not. Numeric capabilities have the value -1 if the capability is missing, and a value at least 0 if it is present. String capabilities (both those with and those without parameters) have the value `NULL` if the capability is missing, and otherwise have type

```
char *
```

and point to a character string containing the capability. The special character codes involving the `\` and `^` characters (such as `\r` for RETURN, or `^A` for CONTROL-A) are translated into the appropriate ASCII characters. Padding information (of the form `$<time>`) and parameter information (beginning with `%`) are left uninterpreted at this

stage. The routine `tputs` interprets padding information, and `tparm` interprets parameter information.

If the program needs to handle only one terminal, the definition `-DSINGLE` can be passed to the C compiler, resulting in static references to capabilities instead of dynamic references. This can result in smaller code, but prevents use of more than one terminal at a time. Very few programs use more than one terminal, so almost all programs can use this flag.

`setupterm(`*term*`,`*filenum*`,`*errret*`)`

This routine is called to initialize a terminal. *term* is the character string representing the name of the terminal being used. *filenum* is the A/UX file descriptor of the terminal being used for output. *errret* is a pointer to an integer, in which a success or failure indication is returned. The values returned can be 1 (all is well), 0 (no such terminal), or -1 (some problem locating the `terminfo` data base).

The value of *term* can be given as 0, which causes the value of `TERM` in the environment to be used. The *errret* pointer also can be given as 0, meaning no error code is wanted. If *errret* is the default, and something goes wrong, `setupterm` prints an appropriate error message and quits, rather than returning. Thus, a simple program can call `setupterm(0,1,0)` and not worry about initialization errors.

If the variable `TERMINFO` is set in the environment to a pathname, `setupterm` checks for a compiled `terminfo` description of the terminal under that path, before checking `/usr/lib/terminfo`. Otherwise, only `/usr/lib/terminfo` is checked.

`setupterm` checks the TTY driver mode bits, using *filenum*, and changes any that might prevent the correct operation of other low-level routines. Currently, the mode that expands tabs into spaces is disabled, because the tab character is sometimes used for different functions by different terminals. (Some terminals use it to move right one space. Others use it to address the cursor to row or column 9.) If the system is expanding tabs, `setupterm` removes the definition of the `tab` and `backtab` functions, making the assumption that because the user is not using hardware tabs, they might not be properly set in the terminal. Other system-dependent changes, such as disabling a virtual terminal driver, can be made here.

As a side effect, `setupterm` initializes the global variable `ttytype`, which is an array of characters, to the value of the list of names for the terminal. This list comes from the beginning of the `terminfo` description.

After the call to `setupterm`, the global variable `cur_term` is set to point to the current structure of terminal capabilities. By calling `setupterm` for each terminal, and saving and restoring `cur_term`, it is possible for a program to use two or more terminals at once.

The mode that turns newlines into "carriage return–line feed" on output is not disabled. Programs that use `cursor_down` or `scroll_forward` should avoid these capabilities if their value is line feed, unless they disable this mode. `setupterm` calls `reset_prog_mode` after any changes it makes.

```
def_prog_mode()
def_shell_mode()
reset_prog_mode()
reset_shell_mode()
```

These routines can be used to change the TTY modes between the two states: shell (the mode they were in before the program was started) and program (the mode needed by the program). `def_prog_mode` saves the current terminal mode as program mode. `setupterm` and `initscr` call `def_shell_mode` automatically. `reset_prog_mode` puts the terminal into program mode, and `reset_shell_mode` puts the terminal into normal mode. A typical calling sequence is for a program to call `initscr` (or `setupterm` if a `terminfo`-level program), then to set the desired program mode by calling routines such as `cbreak` and `noecho`, and then to call `def_prog_mode` to save the current state. Before a shell escape or CONTROL-Z suspension, the program should call `reset_shell_mode`, to restore normal mode for the shell. Then, when the program resumes, it should call `reset_prog_mode`. Also, all programs must call `reset_shell_mode` before they quit. (The higher-level routine `endwin` automatically calls `reset_prog_mode`.)

Normal mode is stored in

```
cur_term->Ottyb,
```

and program mode is in

```
cur_term->Nttyb
```

These structures are both of type `SGTTYB` (which varies depending on the system). Currently the possible types are

```
struct sgttyb
```

(on some other systems) and

```
struct termio
```

(on this version of the A/UX system). `def_prog_mode` should be called to save the current state in `Nttyb`.

`vidputs(`*newmode, putc*`)`
`newmode` is any combination of attributes, defined in `<curses.h>`. `putc` is a `putchar`-like function. The proper string to `put` the terminal in the given video mode is generated. The previous mode is remembered by this routine. The result characters are passed through `putc`.

`vidattr(`*newmode*`)`
The proper string to `put` the terminal in the given video mode is output to `stdout`.

`tparm(`*instring, p1, p2, p3, p4, p5, p6, p7, p8, p9*`)`
`tparm` is used to instantiate a parameterized string. The character string returned has the given parameters applied, and is suitable for `tputs`. Up to nine parameters can be passed, in addition to the parameterized string.

`tputs(`*cp, affcnt, outc*`)`
A string capability, possibly containing padding information, is processed. Enough padding characters to delay for the specified time replace the padding specification, and the resulting string is passed, one character at a time, to the routine *outc*, which should expect one character parameter. (This routine often just calls `putchar`.) *cp* is the capability string. *affcnt* is the number of units affected by the capability, which varies with the particular capability. (For example, the *affcnt* for `insert_line` is the number of lines below the inserted line on the screen—that is, the number of lines that must be moved by the terminal.) *affcnt* is used by the padding information of some terminals as a multiplication factor. If the capability does not have a factor, the value 1 should be passed.

`putp(str)`
This is a convenient function to output a capability with no *affcnt*. The string is output to `putchar` with an *affcnt* of 1. It can be used in simple applications that do not need to process the output of `tputs`.

`delay_output(ms)`
A delay is inserted into the output stream for the given number of milliseconds. The current implementation inserts sufficient pad characters for the delay. This should not be used in place of a high-resolution sleep, but rather for delay effects in the output. Due to buffering in the system, it is unlikely that this call results in the process actually sleeping. Because large numbers of pad characters can be generated, it is recommended that *ms* not exceed 500.

Operation details

These paragraphs describe many of the details of how the `curses` and `terminfo` package operates.

Insert and delete line and character

The algorithm used by `curses` takes into account insert and delete line and character functions, if available, in the terminal. Calling the routine

```
idlok(stdscr, TRUE);
```

enables insert/delete line. By default, `curses` does not use insert/delete line. This was omitted for performance reasons, because there is no speed penalty involved. Rather, experience shows that some programs do not need this facility, and that if `curses` uses insert/delete line, the result on the screen can be visually annoying. Many simple programs using `curses` do not need this, so the default is to avoid insert/delete line. Insert/delete character is always considered.

Additional terminals

`curses` works even if absolute cursor addressing is not possible, as long as the cursor can be moved from any location to any other location. It considers local motions, parameterized motions, home, and carriage return.

`curses` is aimed at full-duplex, alphanumeric, video terminals. No attempt is made to handle half-duplex, synchronous, hard copy, or bit-mapped terminals. Bit-mapped terminals can be handled by programming the bit-mapped terminal to emulate an ordinary alphanumeric terminal. This does not take advantage of the bit-map capabilities, but it is the fundamental nature of `curses` to deal with alphanumeric terminals.

The `curses` package handles terminals with the "magic-cookie glitch" in their video attributes. The term **magic cookie** means that a change in video attributes is implemented by storing a magic cookie in a location on the screen. This cookie takes up a space, preventing an exact implementation of what the programmer wanted. `curses` takes the extra space into account and moves part of the line to the right, as necessary. Advantage is taken of existing spaces, but in some cases this unavoidably results in losing text from the right edge of the screen.

Multiple terminals

Some applications need to display text on more than one terminal, controlled by the same process. Even if the terminals are of different types, `curses` can handle this. All information about the current terminal is kept in a global variable

```
struct screen *SP;
```

Although the screen structure is hidden from the user, the C compiler accepts declarations of variables that are pointers. The user program should declare one screen pointer variable for each terminal it wants to handle. The routine

`struct screen *newterm(`*type*`,`*fd*`)`

sets up a new terminal of the given terminal type, which does output on file descriptor *fd*. A call to `initscr` is essentially

```
newterm(getenv("TERM"),stdout)
```

A program wanting to use more than one terminal should use `newterm` for each terminal and save the value returned as a reference to that terminal.

To switch to a different terminal, call

`set_term(`*term*`)`

The old value of `SP` is returned. The programmer should not assign directly to `SP` because certain other global variables must also be changed.

All `curses` routines always affect the current terminal. To handle several terminals, switch to each one in turn with `set_term`, and then access it. Each terminal must be set up with `newterm`, and closed down with `endwin`.

Video attributes

Video attributes can be displayed in any combination on terminals with this capability. They are treated as an extension of the standout capability, which is still present.

Each character position on the screen has 16 bits of information associated with it. Seven of these bits are the character to be displayed, leaving separate bits for nine video attributes. These bits are used for standout, underline, reverse video, blink, dim, bold, blank, protect, and alternate character set. Standout is taken to be whatever highlighting works best on the terminal, and should be used by any program that does not need specific or combined attributes. Underlining, reverse video, blink, dim, and bold are the usual video attributes. Blank means that the character is displayed as a space, for security reasons. Protect and alternate character set depend on the particular terminal. The use of these last three bits is subject to change and not recommended. Note also that not all terminals implement all attributes—in particular, no current terminal implements both dim and bold.

The routines to use these attributes include

`attrset(`*attrs*`)`	`wattrset(win, `*attrs*`)`
`attron(`*attrs*`)`	`wattron(win, `*attrs*`)`
`attroff(`*attrs*`)`	`wattroff(`*win*`, `*attrs*`)`
`standout()`	`wstandout(`*win*`)`
`standend()`	`wstandend(`*win*`)`

Attributes, if given, can be any combination of

```
A_STANDOUT
A_UNDERLINE
A_REVERSE
A_BLINK
A_DIM
A_BOLD
A_INVIS
A_PROTECT
A_ALTCHARSET
```

These constants, defined in `curses.h`, can be combined with the C language OR operator (`|`) to get multiple attributes.

`attrset(`*attrs*`)`	Sets the current attributes to the given *attrs*
`attron(`*attrs*`)`	Turns on the given *attrs* in addition to any attributes that are already on
`attroff(`*attrs*`)`	Turns off the given *attrs*, without affecting any others

`standout()`
`standend()`
	Equivalent to
	`attron(A_STANDOUT)`
	`attrset(A_NORMAL)`

If the particular terminal does not have the particular attribute or combination requested, `curses` attempts to use some other attribute in its place. If the terminal has no highlighting at all, all attributes are ignored.

Special keys

Many terminals have special keys, such as arrow keys, to erase the screen or insert or delete text, and keys intended for user functions. The particular sequences these terminals send differ from terminal to terminal. `curses` allows the programmer to handle these keys.

A program using special keys should turn on the keypad by calling

```
keypad(stdscr, TRUE)
```

at initialization. This causes special characters to be passed through to the program by the function `getch`. These keys have constants that are listed in the section "Input" earlier in this chapter. They have values starting at 0401, so they should not be stored in a `char` variable, as significant bits will be lost.

A program using special keys should avoid using the ESCAPE key, because most sequences start with escape, creating an ambiguity. `curses` sets a 1-second alarm to deal with this ambiguity, which causes delayed response to the ESCAPE key. It is a good idea to avoid escape in any case, since there is eventually pressure for nearly *any* screen-oriented program to accept arrow-key input.

Scrolling region

There is a programmer-accessible scrolling region. Normally, the scrolling region is set to the entire window, but the calls

`setscrreg(`*top, bot*`)`

`wsetscrreg(`*win, top, bot*`)`

set the scrolling region for `stdscr` or the given window to any combination of top and bottom margins. When scrolling past the bottom margin of the scrolling region, the lines in the region move up one line, destroying the top line of the region. If scrolling is enabled with `scrollok`, scrolling takes place only within that window. Note that the scrolling region is a software feature and only causes a window data structure to scroll. This might or might not translate to use of the hardware scrolling-region feature of a terminal or of insert/delete line; some "intelligent" terminals perform these operations rather than being controlled directly by the software.

`mini-curses`

`curses` copies from the current window to an internal screen image for every call to `refresh`. If the programmer is interested only in screen-output optimization and does not want the windowing or input functions, an interface to the lower-level routines is available. This makes the program somewhat smaller and faster. The interface is a subset of full `curses`, so that conversion between the levels is not necessary to switch from `mini-curses` to full `curses`.

The following functions of `curses` and `terminfo` are available to the user of `mini-curses`:

`addch(`*ch*`)`	`addstr(`*str*`)`	`attroff(`*attrs*`)`
`attron(`*attrs*`)`	`ttrset(`*at*`)`	`clear()`
`erase()`	`initscr`	`move(`*y*`,`*x*`)`
`mvaddch(`*y*`,`*x*`,`*ch*`)`	`mvaddstr(`*y*`,`*x*`,`*str*`)`	`newterm`
`refresh()`	`standend()`	`standout()`

The following functions of `curses` and `terminfo` are not available to the user of `mini-curses`:

`box`	`clrtobot`	`clrtoeol`
`delch`	`deleteln`	`delwin`
`getch`	`getstrs`	`inch`
`insch`	`insertln`	`longname`
`makenew`	`mvdelch`	`mvgetch`
`mvgetstr`	`mvinch`	`mvinsch`
`mvprintw`	`mvscanw`	`mvwaddch`
`mvwaddstr`	`mvwdelch`	`mvwgetch`
`mvwgetstr`	`mvwin`	`mvwinch`
`mvwinsch`	`mvwprintw`	`mvwscanw`
`newwin`	`overlay`	`overwrite`
`printw`	`putp`	`scanw`
`scroll`	`setscrreg`	`subwin`
`touchwin`	`vidattr`	`waddch`
`waddstr`	`wclear`	`wclrtobot`
`wclrtoeol`	`wdelch`	`wdeleteln`
`werase`	`wgetch`	`wgetstr`
`winsch`	`winsertln`	`wmove`
`wprintw`	`wrefresh`	`wscanw`
`wsetscrreg`		

The subset mainly requires the programmer to avoid use of more than the one-window `stdscr`. Thus, all functions beginning with `w` are generally undefined. Certain high-level functions that are convenient but not essential are also not available, including `printw` and `scanw`. Also, the input routine `getch` cannot be used with

mini-curses. Features implemented at a low level, such as use of hardware insert/delete line and video attributes, are available in both versions. Also, mode-setting routines such as crmode and noecho are allowed.

To access mini-curses, add -DMINICURSES to the CFLAGS in the makefile. If routines are requested that are not in the subset, the loader prints error messages such as

```
Undefined:
m_getch
m_waddch
```

to tell you that the routines getch and waddch were used but are not available in the subset. Because the preprocessor is involved in the implementation of mini-curses, the entire program must be recompiled when changing from one version to the other.

TTY-mode functions

In addition to the save/restore routines savetty and resetty, standard routines are available for going into and out of normal TTY mode. These routines are resetterm, which puts the terminal back in the mode it was in when curses was started; fixterm, which undoes the effects of resetterm—that is, restores the "current curses mode"; and saveterm, which saves the current state to be used by fixterm. endwin automatically calls resetterm, and the routine to handle CONTROL-Z (on other systems that have process control) also uses resetterm and fixterm. Programmers should use these routines before and after shell escapes, and also if they write their own routine to handle CONTROL-Z. These routines are also available at the terminfo level.

Typeahead check

If the user types something during an update, the update stops, pending a future update. This is useful when the user hits several keys, each of which causes a good deal of output. For example, in a screen editor, if the user presses the "forward screen" key, which draws the next screenful of text, several times rapidly, rather than drawing several screens of text, the updates are cut short, and only the last screenful is actually displayed. This feature is automatic and cannot be disabled.

getstr

No matter what the setting of the `stty` *echo* is, strings typed in here are echoed at the current cursor location. The user's *erase* and *kill* characters are understood and handled. This makes it unnecessary for an interactive program to deal with erase, kill, and echoing when the user is typing a line of text.

longname

The `longname` function does not need any arguments. It returns a pointer to a static area containing the actual long name of the terminal.

`nodelay` mode

The call

```
nodelay(stdscr, TRUE)
```

puts the terminal in `nodelay` mode. While in this mode, any call to `getch` returns -1 if there is nothing waiting to be read immediately. This is useful for writing programs requiring "real-time" behavior, where the users watch action on the screen and press a key when they want something to happen. For example, the cursor can be moving across the screen, in real time. When it reaches a certain point, the user can press an arrow key to change direction at that point.

Portability

Several useful routines are provided to improve portability. The implementation of these routines is different from system to system, and the differences can be isolated from the user program by including them in `curses`.

`erasechar()` Returns the character that erases one character.

`killchar()` Returns the character that kills the entire input line.

`baudrate()` Returns the current baud rate as an integer. (For example, at 9600 baud, the integer 9600 is returned, not the value `B9600` from `<sgtty.h>`.)

`flushinp()` Causes all typeahead to be thrown away.

Example program: `scatter`

```
/*
 * scatter: this program takes the first
 * screenful of lines from the standard
 * input and displays them on the
 * VDU screen, in a random manner.
 */
#include <curses.h>
#define MAXLINES 120
#define MAXCOLS 160
char s[MAXLINES][MAXCOLS];      /* Screen array */
main()
{
   register int row,col;
   register char c;
   int char_count=0;
   long t;
   char buf[BUFSIZ];
   initscr();
   for(row=0;row<MAXLINES;row++)
      for(col=0;col<MAXCOLS;col++)
         s[row][col]=' ';
```

(continued)➡

```
row = 0;
/* Read screen in */
while( (c=getchar()) != EOF && row < LINES ) {
   if(c != '\n') {
      /* Place char in screen array */
      s[row][col++] = c;
      if(c != ' ')
         char_count++;
   } else {
       col=0;
       row++;
   }
}
time(&t);       /* Seed random number generator */
  srand((int)(t&0177777L));
  while(char_count) {
    row=rand() % LINES;
    col=(rand()>>2) % COLS;
    if(s[row][col] != ' ')
    {
        move(row, col);
        addch(s[row][col]);
        s[row][col]=EOF;
        char_count--;
        refresh();
     }
  }
  endwin();
  exit(0);
}
```

Example program: `show`

```
/*
 * The show program pages through
 * a file, showing one full screen each
 * time the user presses the space bar
 */
#include <curses.h>
#include <signal.h>
main(argc, argv)
int argc;
char *argv[];
{
   FILE *fp;
   char linebuf[BUFSIZ];
   int line;
   void done(), perror(), exit();
   if(argc != 2)
   {
      fprintf(stderr,"usage: %s file\n", argv[0]);
      exit(1);
   }
   if((fp=fopen(argv[1],"r")) == NULL)
   {
      perror(argv[1]);
      exit(2);
   }
   signal(SIGINT, done);
   initscr();
   noecho();
   cbreak();
   nonl();
   idlok(stdscr, TRUE);
```

(continued)➡

```
    while(1)
    {
        move(0,0);
        for(line=0; line<LINES; line++)
        {
        if(fgets(linebuf, sizeof linebuf, fp) == NULL)
        {
            clrtobot();
            done();
        }
        move(line, 0);
        print("%s", linebuf);
        }
        refresh();
        if(getch() == 'q')
        done();
    }
}
void
done()
{
    move(LINES-1, 0);
    clrtoeol();
    refresh();
    endwin();

}
```

Example program: `highlight`

```
/*
 * highlight: a program to turn \U, \B, and
 * \N sequences into highlighted
 * output, allowing words to be
 * displayed underlined or in bold.
 */
#include <curses.h>
main(argc, argv)
char **argv;
{
   FILE *fp;
   int c, c2;
   if (argc != 2) {
      fprintf(stderr, "Usage: highlight file\n");
      exit(1);
   }
   fp = fopen(argv[1], "r");
   if (fp == NULL) {
      perror(argv[1]);
      exit(2);
   }
   initscr();
   scrollok(stdscr, TRUE);
```

(continued)➡

```
    for (;;) {
        c = getc(fp);
        if (c == EOF)
            break;
        if (c == '\\') {
            c2 = getc(fp);
        switch (c2) {
        case 'B':
            attrset(A_BOLD);
            continue;
        case 'U':
            attrset(A_UNDERLINE);
            continue;
        case 'N':
            attrset(0);
            continue;
        }
        addch(c);
        addch(c2);
        }
        else
        addch(c);
    }
    fclose(fp);
    refresh();
    endwin();
    exit(0);
}
```

Example program: `window`

```
/*
 * This program shows the use of multiple windows.
 * The main display is kept in stdscr.
 * When the user temporarily wants to put
 * something else on the screen,
 * a new window is created covering
 * part of the screen.
 */
#include <curses.h>
WINDOW *cmdwin;
main()
{
   int i, c;
   char buf[120];
   initscr();
   nonl();
   noecho();
   cbreak();
   /* top 3 lines */
   cmdwin = newwin(3, COLS, 0, 0);
   for (i=0; i<LINES; i++)
      mvprintw(i, 0, "This is line %d of stdscr", i);
```

(continued)➡

```
    for (;;) {
        refresh();
        c = getch();
        switch (c) {
        case 'c':       /* Enter command from keyboard */
          werase(cmdwin);
          wprintw(cmdwin, "Enter command:");
          wmove(cmdwin, 2, 0);
          for (i=0; i<COLS; i++)
          waddch(cmdwin, '-');
          wmove(cmdwin, 1, 0);
          touchwin(cmdwin);
          wrefresh(cmdwin);
          wgetstr(cmdwin, buf);
          touchwin(stdscr);
          /*
          * The command is now in buf.
          * It should be processed here.
          */
          break;
         case 'q':
          endwin();
          exit(0);
         }
    }
}
```

Example program: `two`

```
/*
 * The two program pages through a file,
 * showing one page to the first terminal and
 * the next page to the second terminal.
 * It then waits for a space to be typed on
 * either terminal, and shows the next
 * page to the terminal typing the space.
 */
#include <curses.h>
#include <signal.h>
struct screen *me, *you;
struct screen *set_term();
FILE *fp, *fpyou;
char linebuf[512];
main(argc, argv)
char **argv;
{
    int done();
    int c;
    if (argc != 4) {
        fprintf(stderr,
        "Usage: two othertty otherttytype inputfile\n");
        exit(1);
    }
    fp = fopen(argv[3], "r");
    fpyou = fopen(argv[1], "w+");
    signal(SIGINT, done);
    /* die gracefully */
```

(continued)➡

```
me = newterm(getenv("TERM"), stdout);
/* initialize my tty */
you = newterm(argv[2], fpyou);
/* Initialize his terminal */
/* Set modes for my terminal */
set_term(me);
noecho();                /* turn off tty echo */
cbreak();                /* enter cbreak mode */
nonl();                  /* Allow linefeed */
nodelay(stdscr,TRUE);    /* No hang on input */
/* Set modes for other terminal */;
set_term(you)
noecho();
cbreak();
nonl();
nodelay(stdscr,TRUE);
/* Dump first screenful on my terminal */
dump_page(me);
/* Dump second screenful on his terminal */
dump_page(you);
/* for each screenful */
for (;;) {
   set_term(me);
   c = getch();
   /* wait for user to read it */)
   if (c == 'q'
   done();
   if (c == ' ')
   dump_page(me);
```

```
        set_term(you);
        c = getch();
        /* wait for user to read it */
        if (c == 'q')
        done();
        if (c == ' ')
        dump_page(you);
        sleep(1);
    }
}
dump_page(term)
struct screen *term;
{
    int line;
    set_term(term);
    move(0, 0);
    for (line=0; line<LINES-1; line++) {
        if (fgets(linebuf, sizeof linebuf, fd) == NULL)
        {
        clrtobot();
        done();
        }
        mvprintw(line, 0, "%s", linebuf);
    }
    standout();
    mvprintw(LINES-1, 0, "--More--");
    standend();
    refresh();                  /* sync screen */
}
```

(continued)➡

```
/*
 * Clean up and exit.
 */
done()
{
    /* Clean up first terminal */
    set_term(you);
    move(LINES-1,0);            /* to lower left corner */
    clrtoeol();                 /* clear bottom line */
    refresh();                  /* flush out everything */
    endwin();                   /* curses cleanup */

    /* Clean up second terminal */
    set_term(me);
    move(LINES-1,0);            /* to lower left corner */
    clrtoeol();                 /* clear bottom line */
    refresh();                  /* flush out everything */
    endwin();                   /* curses cleanup */

    exit(0);
}
```

Example program: `termhl`

```
/*
 * A terminfo-level version of highlight.
 */
#include <curses.h>
#include <term.h>
int ulmode = 0;          /* Currently underlining */
main(argc, argv)
char **argv;
{
   FILE *fp;
   int c, c2;
   int outch();
   if (argc > 2) {
      fprintf(stderr, "Usage: termhl [file]\n");
      exit(1);
   }
   if (argc == 2) {
      fp = fopen(argv[1], "r");
      if (fp == NULL) {
      perror(argv[1]);
      exit(2);
      }
   } else {
      fp = stdin;
   }
   setupterm(0, 1, 0);
```

(continued)➡

```
    for (;;) {
        c = getc(fp);
        if (c == EOF)
        break;
        if (c == '\\') {
        c2 = getc(fp);
        switch (c2) {
        case 'B':
            tputs(enter_bold_mode, 1, outch);
            continue;
        case 'U':
            tputs(enter_underline_mode, 1, outch);
            ulmode = 1;
            continue;
        case 'N':
            tputs(exit_attribute_mode, 1, outch);
            ulmode = 0;
            continue;
        }
        putch(c);
        putch(c2);
        }
        else
        putch(c);
    }
    fclose(fp);
    fflush(stdout);
    resetterm();
    exit(0);
}
```

```
/*
 * This function is like putchar,
 * but it checks for underlining.
 */
putch(c)
int c;
{
   outch(c);
   if (ulmode && underline_char) {
      outch('\b');
      tputs(underline_char, 1, outch);
   }
}
/*
 * Outchar is a function version
 * of putchar that can be passed to
 * tputs as a routine to call.
 */
outch(c)
int c;
{
   putchar(c);
}
```

Example program: `editor`

```
/*
 * editor: A screen-oriented editor. The user
 * interface is similar to a subset of vi.
 * The buffer is kept in stdscr itself to simplify
 * the program.
 */

#include <curses.h>

#define CTRL(c) ('c' & 037)

main(argc, argv)
int argc;
char **argv;
{
    int i, n, l;
    int c;
    FILE *fp;

    if (argc != 2) {
        fprintf(stderr, "Usage: edit file\n");
        exit(1);
    }

    fp = fopen(argv[1], "r");
    if (fp == NULL) {
        perror(argv[1]);
        exit(2);
    }

    initscr();
    cbreak();
    nonl();
    noecho();
    idlok(stdscr, TRUE);
    keypad(stdscr, TRUE);
```

```
    /* Read in the file */
    while ((c = getc(fp)) != EOF)
        addch(c);
    fclose(fp);
    move(0,0);
    refresh();
    edit();
    /* Write out the file */
    fp = fopen(argv[1], "w");
    for (l=0; l<23; l++) {
        n = len(l);
        for (i=0; i<n; i++)
        putc(mvinch(l, i), fp);
        putc('\n', fp);
    }
    fclose(fp);
    endwin();
    exit(0);
}
len(lineno)
int lineno;
{
    int linelen = COLS-1;
    while (linelen >=0
          && mvinch(lineno,linelen) == ' ') linelen--;
    return linelen + 1;
}
/* Global value of current cursor position */
int row, col;
edit()
{
    int c;
```

(continued)➡

```
for (;;) {
   move(row, col);
   refresh();
   c = getch();
   switch (c) {         /* Editor commands */
   /* hjkl and arrow keys: move cursor */
   /* in direction indicated */
   case 'h':
   case KEY_LEFT:
   if (col > 0)
      col--;
   break;
   case 'j':
   case KEY_DOWN:
   if (row < LINES-1)
      row++;
   break;
   case 'k':
   case KEY_UP:
   if (row > 0)
   row--;
   break;
   case 'l':
   case KEY_RIGHT:
   if (col < COLS-1)
      col++;
   break;
   /* i: enter input mode */
   case KEY_IC:
   case 'i':
   input();
   break;
```

```
/* x: delete current character */
case KEY_DC:
case 'x':
delch();
break;
/* o: open up a new line and enter input mode */
case KEY_IL:
case 'o':
move(++row, col=0);
insertln();
input();
break;
/* d: delete current line */
case KEY_DL:
case 'd':
deleteln();
break;
/* ^L: redraw screen */
case KEY_CLEAR:
case CTRL(L):
clearok(curscr);
refresh();
break;
/* w: write and quit */
case 'w':
return;
```

(continued)➡

```
        /* q: quit without writing */
        case 'q':
        endwin();
        exit(1);
        default:
        flash();
        break;
        }
    }
}
/*
 * Insert mode: accept characters and insert them.
 * End with ^D or EIC
 */
input()
{
    int c;
    standout();
    mvaddstr(LINES-1, COLS-20, "INPUT MODE");
    standend();
    move(row, col);
    refresh();
    for (;;) {
        c = getch();
        if (c == CTRL(D) || c == KEY_EIC)
        break;
        insch(c);
        move(row, ++col);
        refresh();
    }
```

```
    move(LINES-1, COLS-20);
    clrtoeol();
    move(row, col);
    refresh();
}
```

13 Commando

Introduction / 13-2
The Commando script language / 13-5
Creating Commando dialogs / 13-30
Dialog design guidelines / 13-32

This chapter explains how you can write Commando dialog scripts to provide a Macintosh front end for your UNIX applications.

Commando lets you create CommandShell command lines by selecting controls within Macintosh dialog boxes. Controls direct the placement of options on the command line. By selecting a particular control, a specific option can be placed on the command line. The command lines thus constructed are placed in a CommandShell window for execution or are optionally executed in a subshell.

This chapter begins with a general discussion of dialog boxes; readers who are familiar with this subject might want to turn directly to the section "Commando Dialog Boxes."

Introduction

The Macintosh computer provides you with visual cues when you communicate with an application, among them the **controls** used in dialog boxes. Controls allow you to change the way an application functions; when a particular control is used it can place a specific option on the command line. The use of dialog boxes provides a consistency of interface across applications that decreases learning times for new applications and increases retention times for completed tasks. By implementing this interface on UNIX applications already running on A/UX, programmers and developers can increase the effectiveness of users working with the application.

Users who are relatively new to command-line interfaces often do not take the time necessary to learn all the intricacies needed to make full use of the features of a program. Further, they are often frustrated in their attempts to use applications because it is not obvious what options are available, or what the application does if they enter a given option. This is where Commando can help. Because Commando translates between visual controls and command-line options, users can see at a glance what an application can do and know what options are available. Further, Commando includes a context-sensitive help feature, so users receive an explanation of the effect of each control as they click it. Programmers can save time because they do not have to explain the workings of the application time and time again.

Commando lets you create command lines using the controls within Macintosh dialog boxes. This makes invocation of even complex commands much easier, since users have feedback on what the command can do before they execute it. This also benefits occasional users of UNIX, because it frees them from having to memorize the options or arguments associated with various commands. Even UNIX experts appreciate this feature, since few have learned all the options of the more than 500 UNIX commands.

The contents of Commando dialog boxes are specified in *dialog scripts* written according to the Commando script language, which is discussed in detail later in this chapter. Much of the work of laying out the dialog boxes, including automatic vertical spacing, is done for you. This leaves you free to concentrate on the logical presentation order of the controls.

The steps you typically follow to create a Commando dialog are quite simple:

1. Copy a Commando dialog script from an existing command having similar controls. Scripts for all the Commando dialogs are kept in directories in `/mac/lib/cmdo`.
2. Modify the script to reflect the controls for your application or utility.
3. Test and debug the script.
4. Make sure the script has read-only permission.
5. Move the script to the appropriate directory in `/mac/lib/cmdo`.

These steps are described in detail later in this chapter.

Macintosh dialog boxes

Dialog boxes provide the user with several visual cues (see Figure 13-1). The use of dialog boxes is governed by several conventions:

- Checkboxes allow users to select an option individually; these are the default controls in Commando.
- Radio buttons allow users to select mutually exclusive options.
- Text boxes allow (or require) users to enter additional information.
- Buttons either allow users to select files or lead to further dialog boxes.
- Controls that cannot be selected are dimmed.

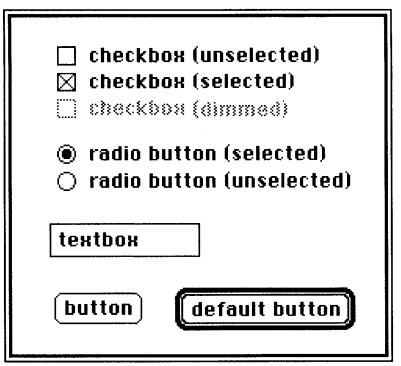

Figure 13-1 Schematic dialog box

Commando dialog boxes

All Commando dialog boxes have similar structures, though the controls for the command they represent are different. Figure 13-2 shows a representative dialog box for a UNIX command. Each dialog shows the current command line being built, a box of Help information, and buttons to send or cancel the command. Each screen also has an area to select among the various options of the command. Each dialog box can have multiple controls, allowing command lines to be arbitrarily complex. Further, each command can have several nested dialog boxes; in Figure 13-2, the "File type," "Fonts," and "More options" buttons each lead to a nested subdialog.

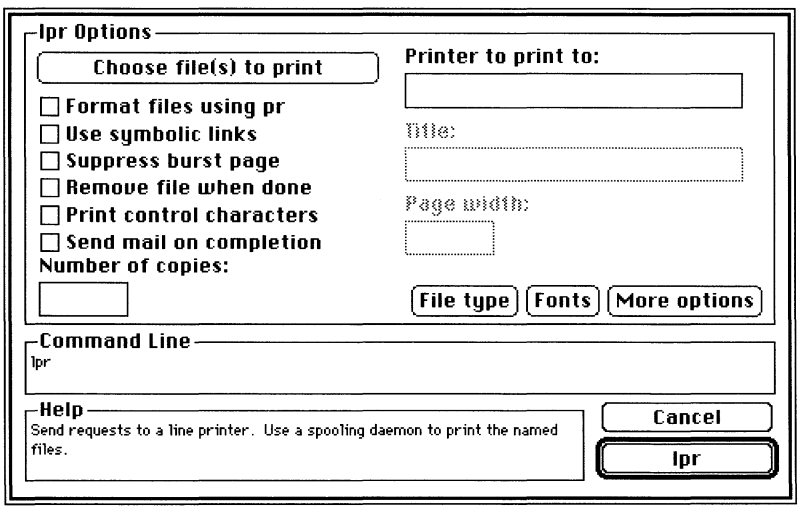

Figure 13-2 Commando dialog box for the UNIX command `lpr`

Controls can be set up to enable other controls. In Figure 13-2, the title and page width controls are disabled because they are used only when the option "Format files using pr" is selected. Since this control hasn't been selected, the title and page width controls are inaccessible. Examples of this enabling feature are shown later in this chapter in the section "Commando Script Language."

Figure 13-3 shows another dialog box, this one for the `tar` command. The Operation controls, which control whether the program reads from or writes to the backup media, are mutually exclusive and thus are implemented as radio buttons. The buttons giving access to dialog boxes containing further controls are enabled only when their corresponding radio button has been selected.

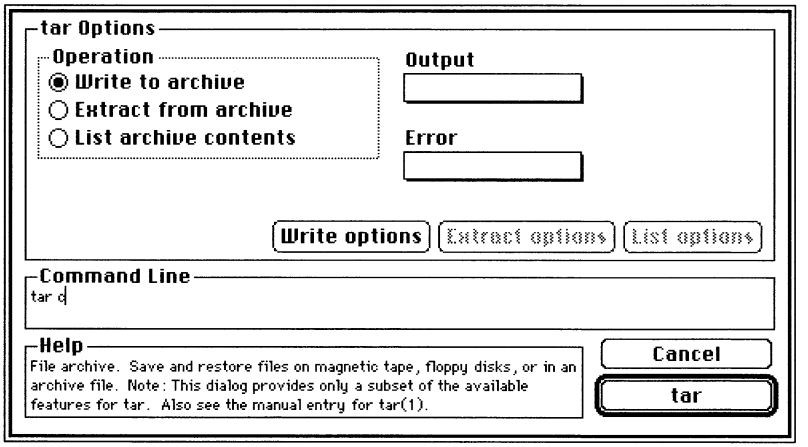

Figure 13-3 Commando dialog box for the UNIX command `tar`

The Commando script language

The Commando script language helps you to create well-designed Commando dialog boxes quickly. Commando scripts allow users to start an application by double-clicking an icon or by invoking the application dialog script from the command line. The resulting dialog boxes allow the user to select various flags and options, then pass the command line to CommandShell for execution. By using dialog boxes developed through Commando, you can give your applications the front end of a Macintosh application without changing the code of your UNIX application.

Dialog box layout

All Commando dialogs have several aspects of their layout in common: All have labeled Options, Command Line, and Help boxes (see Figure 13-4). All have button controls in the lower-right corner of the box, allowing you to cancel the displayed dialog box or (by default) complete your selections and send the command line to CommandShell.

The Option box of the dialog is laid out in rows and columns. There can be several rows within a given box. You can have multiple columns within a row, and multiple rows within a column.

Within a column or row are various command controls. Buttons, checkboxes, text boxes, and radio buttons define how the command line is built. Additional outline boxes can be added to group similar functions visually. Optional definitions might require or enable controls. Buttons leading to other dialog boxes can be included.

Figure 13-4 shows an Option box layout having two rows, a and b, enclosed within column 1, and the two columns, 1 and 2, enclosed within one large row, A. The various rows and columns are indicated by rectangles and names (in operation, Commando does not draw these rectangles or insert the names unless you specifically direct it to). In this simple layout example, no programmer-defined controls are shown.

Just as all Commando dialogs have some structures in common, all scripts have some structures in common. The beginning of the script always defines the name of the command, in this case "sample," by using the keyword `command name`. The name of the command appears in the default invocation button, in the Command Line box, and at the top left of the Option box (see Figure 13-4). Next, the keyword `help` defines the message shown in the Help box when you are not clicking a specific control. The maximum length of a help message varies with the length of the command name, but roughly 200 characters can be included.

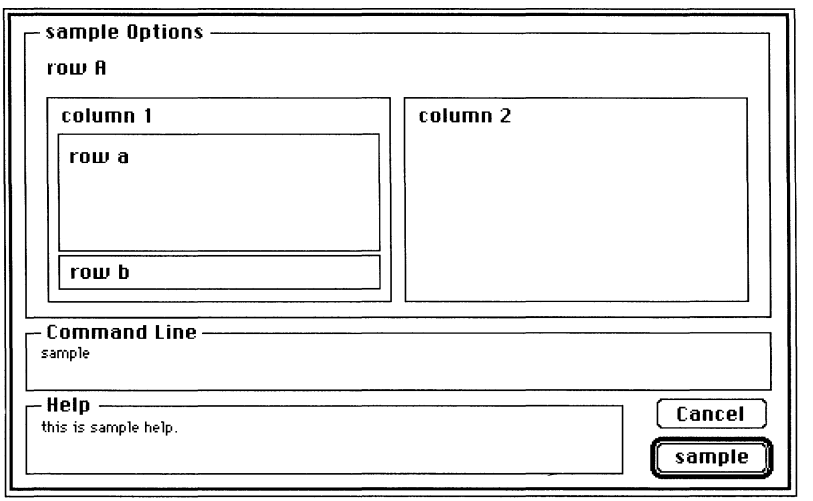

Figure 13-4 Dialog box layout example

The remainder of the script defines rows and columns of controls. Scripts reflect the structure displayed. If you want multiple columns within a row, column definitions are nested within the row definition. Each definition for a particular row or column is enclosed by braces. Row definitions begin with the keyword `row`, and column definitions with the keyword `column`. The braces might enclose other layout or control keywords, further affecting the appearance of the dialog box. (Commando automatically adjusts the vertical size to include the defined controls.) Each keyword begins on its own line in a dialog script. (A complete list of keywords is given in the section "Keywords.")

The script shown in Listing 13-1 reflects the structure displayed in Figure 13-4. Definitions for the innermost rows, a and b, are nested within column 1. The definitions for columns 1 and 2 are nested within row A. This simple example does not include any control specifiers; they would be enclosed by the braces between the beginning and end of the definition of a column or row.

◆ **Note** Both within "real" dialog scripts and in the following examples comments are bracketed by slashes and asterisks: `/* this is a comment */`. Comments are used in the following examples to point out specific features of dialog scripts. Comments can be only one line long. ◆

Listing 13-1 Dialog box layout example script

```
command name "sample"
help "this is sample help."
row {                    /* this begins row A */
  column {                  /* this begins column 1 */
        row {                  /* this begins row a */
            }                  /* this ends row a */
        row {                  /* this begins row b */
            }                  /* this ends row b */
         }                 /* this ends column 1 */
  column {                 /* this begins column 2 */
         }                 /* this ends column 2 */
    }                    /* this ends row A */
```

Layout examples

This section contains examples of the layout of controls. The controls themselves are not discussed in depth; they are discussed later in the chapter in the section "Control Examples."

Single-row example

Figure 13-5 shows a trivial example containing only two programmer-defined controls. In this example they are both checkboxes. The dialog script that produced this dialog is shown in Listing 13-2.

◆ **Note** In the following examples a sample Commando dialog box is shown first, and the dialog script that produced it is shown next. ◆

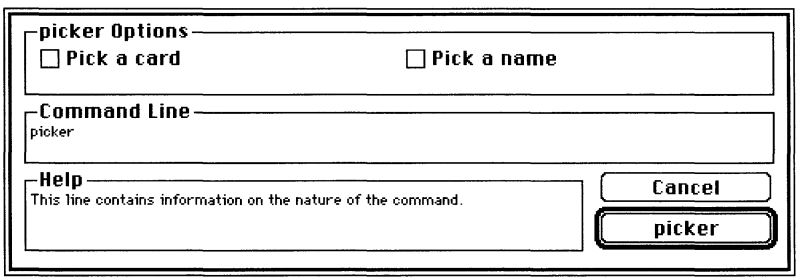

Figure 13-5 Single-row dialog box

Listing 13-2 Single-row dialog script

```
command name "picker"                    /* command name */

help  "This line contains information "  /* help message */

      "on the nature of the command."

row {                                    /* begin only row */

    option name "Pick a card"            /* first control */

        prefix "-p"
        help "[-p] randomly pick a card."

    option name "Pick a name"            /* second control */
```

```
        prefix "-n"
        help "[-n] randomly pick a name."
    }                                        /* end of row */
```

A line-by-line dissection of this script shows several general points worth noting:

- The first line,

  `command name "picker"`

  defines the command name. As mentioned previously, it is automatically put on the command line being built (in the Command Line box of Figure 13-5), in the command button (at the lower right of Figure 13-5), and in the Options label at the top of the dialog box.

- The second and third lines,

  ```
  help  "This line contains information "
        "on the nature of the command."
  ```

  define the help message for the command displayed in the bottom Help box. The message can span several lines, which are concatenated when the dialog is constructed. The help message for the command is displayed whenever the mouse button is up.

- The fourth line,

  `row {`

  specifies construction of a row. Controls between this point and the closing brace (on the last line) are all placed on the same row.

- On the fifth line,

  `option name "Pick a card"`

  the `option name` for the first control defines how the control is to be labeled.

- On the sixth line,

  `prefix "-p"`

  the `prefix` line defines what characters are placed on the command line being built when this control is selected.

- On the seventh line,

  `help "[-p] randomly pick a card."`

  the `help` line following an `option name` keyword defines what appears in the Help box when the pointer is positioned over this control and the mouse button is down.

- The eighth through tenth lines,

```
option name "Pick a name"
   prefix "-n"
   help "[-n] randomly pick a name."
```

  specify another control; each control consists of at least the option name, prefix, and a help message.

- The eleventh line,

```
}
```

  has the closing brace for the first row.

Commando automatically divides a row into columns to space the controls. In Figure 13-7 there are two controls, so two columns are used for spacing. This spacing can affect the length you choose for control names.

Multiple-row example

Figure 13-6 shows a dialog box with different rows having different numbers of controls. The first row contains the three controls: "Pick a card," "Pick a name," and "Pick a spot." The second row contains the two pop-up menus Output and Error.

This example shows what the dialog box looks like if the pointer is positioned on the "Pick a card" option and the mouse button is down. The control shows that it is selected (there is an X in the checkbox), the `-p` prefix shows in the Command Line box, and the Help box displays the message associated with that option. When the mouse button is released, the control remains selected and the prefix remains in the command line being built, but the help message reverts to the message for the command itself.

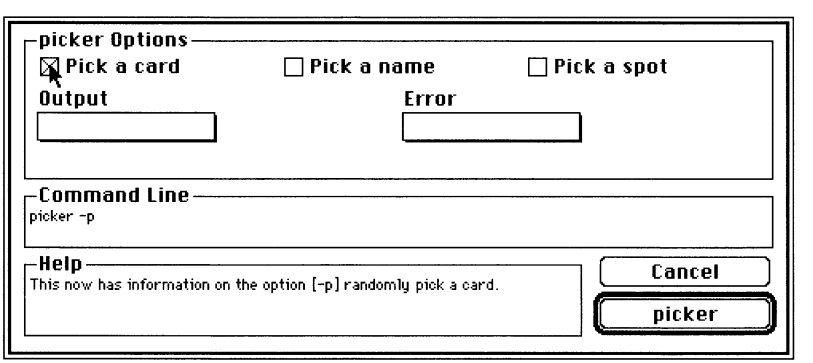

Figure 13-6 Multiple-row dialog box

Figure 13-6 and Listing 13-3 show a new point: The type of control displayed in the dialog changes when the keyword within an `option name` section is changed. Note that the options within the bold area of Listing 13-3 use the keywords `outpopup` and `errpopup`. These keywords create different kinds of controls from those created by the default checkbox. The various types of controls are covered in depth in the section "Control Examples." These examples again demonstrate the automatic building of columns within rows. The first row has three controls and is displayed in three columns, while the second row has two controls and is displayed in two columns. The vertical spacing is again adjusted automatically to allow room for the controls.

Listing 13-3 Multiple-row dialog script

```
command name "picker"                          /* command name */

help  "This line contains information "        /* help message */
      "on the nature of the command."
row {                                        /* start first row */
    option name "Pick a card"                  /* first control */
        prefix "-p"
        help "This now has information on the option "
             "[-p] randomly pick a card."
    option name "Pick a name"                 /* second control */
        prefix "-n"
        help ""This now has information on the option "
             "[-n] randomly pick a name."
    option name "Pick a spot"                  /* third control */
        prefix "-s"
        help ""This now has information on the option "
             "[-s] randomly pick a spot."
    }                                          /* end first row */
row {                                       /* start second row */
    option name "output"          /* first control */
        outpopup
    option name "error"           /* second control */
        errpopup
    }                                        /* end second row */
```

Column example

The following examples (Figure 13-7 and Listing 13-4) demonstrate the explicit definition of a column. Multiple columns can be defined within a row, with the horizontal spacing divided equally by the defined number of columns. Multiple columns containing different numbers of controls can be contained within the same row. Commando automatically adjusts the vertical height of the dialog box based on the number of controls in a particular column (within limits, of course).

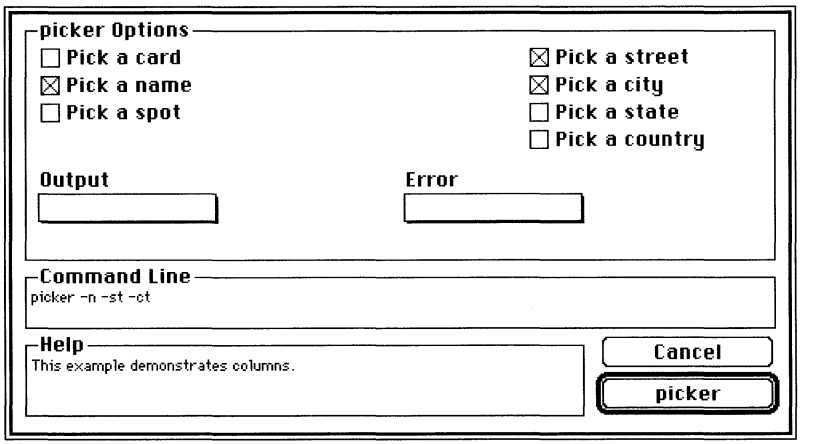

Figure 13-7 Multiple-column dialog box

In the first bold area of Listing 13-4 the keyword `column` is used within the first row to put all three controls in the same column. The plain area between the bold areas contains a dummy column, one with nothing between its braces; it is used to create the blank column. The lower bold area starts another column specification, this time putting four controls in the column. Commando again takes care of adjusting the vertical spacing of the dialog box.

Listing 13-4 Multiple-column dialog script

```
command name "picker"                          /* command name */
help  "This example demonstrates columns."     /* help message */

row {                                  /* start first row */

    column {                       /* start first column */
```

```
        option name "Pick a card"           /* first control */
            prefix "-p"
            help "[-p] randomly pick a card."
        option name "Pick a name"          /* second control */
            prefix "-n"
            help "[-n] randomly pick a name."
        option name "Pick a spot"          /* third control */
            prefix "-s"
            help "[-s] randomly pick a spot."
           }                              /* end first column */
    column {}                  /* dummy second column for spacing */
    column {                        /* start third column */
        option name "Pick a street"        /* first control */
            prefix "-st"
            help "[-s] randomly pick a street."
        option name "Pick a city"          /* second control */
            prefix "-ct"
            help "[-n] randomly pick a city."
        option name "Pick a state"         /* third control */
            prefix "-sta"
            help "[-n] randomly pick a state."
        option name "Pick a country"       /* fourth control */
            prefix "-c"
            help "[-s] randomly pick a country."
           }                              /* end third column */
    }                                    /* end first row */
row {                                    /* start second row */
    option name "output"                     /* first control */
        outpopup
    option name "error"                      /* second control */
        errpopup
    }                                    /* end second row */
```

Nested dialog box example

Figure 13-8 shows the next step in changing the structure, the addition of a button leading to a further dialog box (nested dialog boxes are also referred to as *subdialogs*). Here, a new dialog named *Redirection* was added to the first dialog box. Clicking the Redirection button leads to the subdialog, shown as the second dialog of Figure 13-8.

Note that the second dialog box (the lower box shown in Figure 13-8) has a Continue button that returns the user to the first dialog box. Multiple nested dialog boxes can be specified (see Figure 13-2, "Commando Dialog Box for the UNIX Command `lpr`").

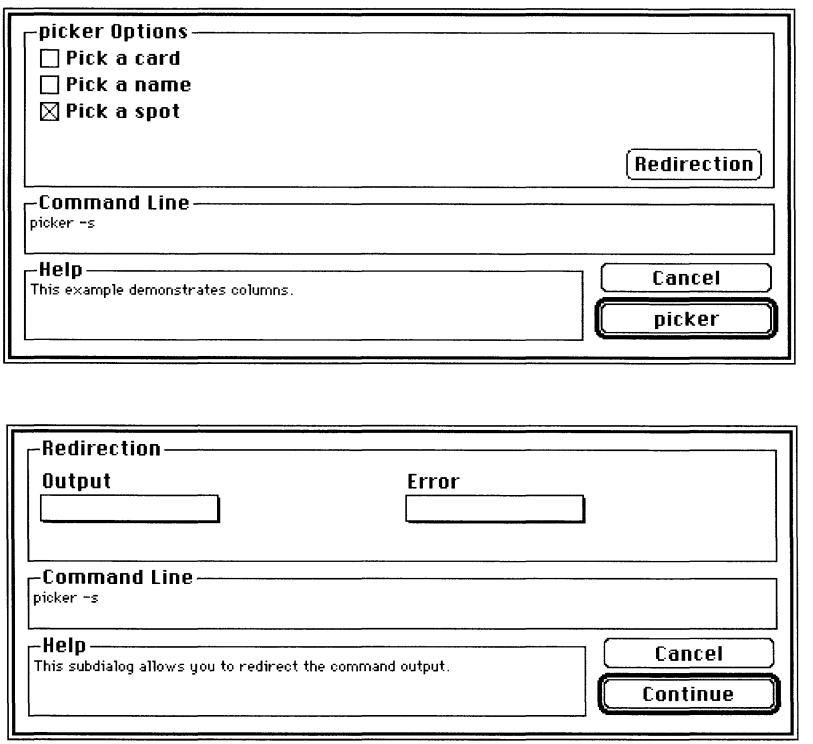

Figure 13-8 Further dialog example

The bold area of Listing 13-5 shows the addition of the button named Redirection, which leads to a further dialog box. Buttons to access further dialogs are automatically sized to hold the button name.

Listing 13-5 Further dialog script

```
command name "picker"                          /* command name */
help  "This example demonstrates columns."  /* help message */
row {                                       /* start first row */
    column {                                   /* start first column */
        option name "Pick a card"
            prefix "-p"
            help "[-p] randomly pick a card."
        option name "Pick a name"
            prefix "-n"
            help "[-n] randomly pick a name."
        option name "Pick a spot"
            prefix "-s"
            help "[-s] randomly pick a spot."
           }                                  /* end first column */
    }                                         /* end first row */
dialog name "redirection"                   /* start second dialog */
help  "This subdialog allows you to "  /* help message */
      "redirect the command output."
row {                                       /* start first row */
    option name "output"
        outpopup
    option name "error"
        errpopup
    }                                       /* end first row */
```

Control examples

Places in a dialog where the user can make a choice are called *controls*. These include checkboxes, radio buttons, text boxes, and buttons. Keyword specifiers define all controls available from the dialog box. The type of control is usually specified (a checkbox is the default). In addition, enabling and requirement dependencies can be defined (see the section "Dependencies"). An enabling dependency makes access to a particular control dependent on the state of some other control. A requirement dependency forces the user to select a control before the command line can be sent to the shell. The various types of controls are discussed in the following sections.

Checkbox

The checkbox is the default control type; it is a square box that the user selects or deselects by clicking it. The user selects each checkbox individually. Figure 13-9 shows examples of this type of control.

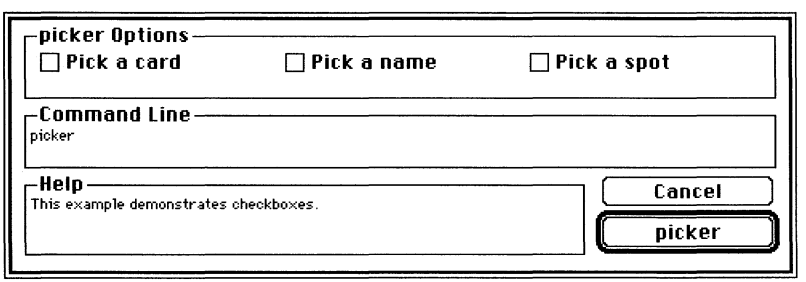

Figure 13-9 Checkbox dialog example

Listing 13-6 shows the script that produced the checkbox dialog example; the bold area contains a representative checkbox definition. You define the `option name`, `prefix`, and `help` specifiers for every checkbox. Place the text following each of these keywords between double quotation marks. The maximum number of checkboxes in a column is ten.

Listing 13-6 Checkbox example script

```
command name "picker"

help "This example demonstrates "
     "checkboxes."
```

```
row {
    option name "Pick a card"
        prefix "-p"
        help "[-p] randomly pick a card."
    option name "Pick a name"
        prefix "-n"
        help "[-n] randomly pick a name."
    option name "Pick a spot"
        prefix "-s"
        help "[-s] randomly pick a spot."
    }
```

Radio buttons

Radio buttons are similar to checkboxes but provide users with mutually exclusive controls; an example is shown in Figure 13-10. Users see associated radio buttons aligned in columns; a box, referred to as a *named box,* usually surrounds radio buttons to visually indicate that they are related. Commando automatically selects the first radio button in a set for the user.

Due to an intentional layout error in the example, one of the labels in Figure 13-10 is too long and has extended outside the named box. This example demonstrates that you must choose control labels that fit within their column.

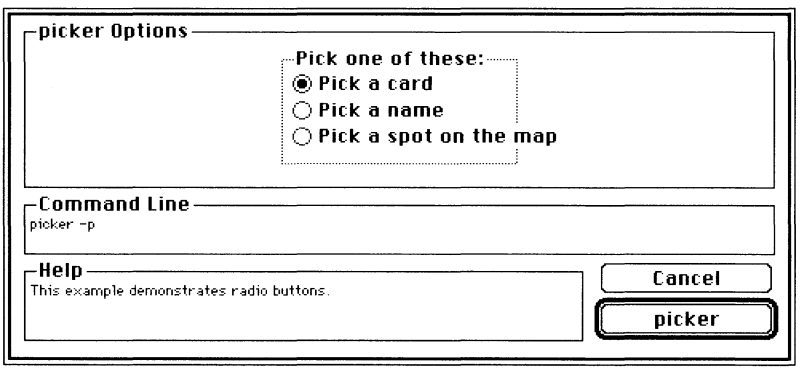

Figure 13-10 Radio button dialog example

The bold area of Listing 13-7 shows a definition for a set of radio buttons. The definition starts with the keyword `radio buttons` and encloses the set of individual controls in braces. Specify the keywords `option name`, `prefix`, and `help` for each button. Use a box to visually indicate the grouping of the radio buttons. To do this, use the keyword `name` or the keyword `box` within the radio button definition. The keyword `name` gives you a labeled box (as shown in Figure 13-10 and Listing 13-7); the keyword `box` creates a simple outline box. Commando automatically aligns radio buttons into columns. A maximum of seven radio buttons can be grouped in a set. By default, Commando selects the first radio button in a set, so make the first control the one that the user most often chooses.

Listing 13-7 Radio button example script

```
command name "picker"

help "This example demonstrates "
     "radio buttons."
row {
    column {}                        /* dummy column for spacing */

    radio buttons {        /* specify a set of radio buttons */
    name "Pick one of these:" /* create a named grouping box */
        option name "Pick a card"
            prefix "-p"
            help "[-p] randomly pick a card."
        option name "Pick a name"
            prefix "-n"
            help "[-n] randomly pick a name."
        option name "Pick a spot on the map"
            prefix "-s"
            help "[-s] randomly pick a spot."
              }

    column {}                        /* dummy column for spacing */
    }
```

Text boxes

A text box allows the user to enter text to be used in the command arguments. Text boxes are the width of the current column. Figure 13-11 shows an example of the use of these text input types. Note that when an input string contains blanks, Commando automatically encloses the string in single quotation marks to avoid confusing the shell. (For example, see the name Wally Eldridge shown in the Command Line box.)

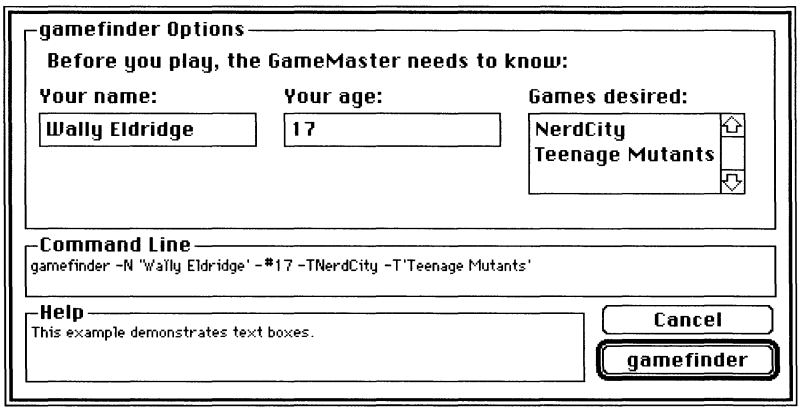

Figure 13-11 Text box dialog example

Listing 13-8 shows the script used to create the dialog in Figure 13-11. As with a checkbox, specify the keywords `option name`, `prefix`, and `help` for each text box. Use one of the keywords `string` or `stringlist` to indicate the type of data to be input by the user. The keyword `string` allows entry of a single line of text, while `stringlist` allows entry of several lines, each of which is prefaced on the command line with the defined `prefix`. Text boxes are the width of the current column. You can put a maximum of three `string` controls or two `stringlist` controls in a column.

As mentioned previously, when an input string contains blanks, Commando automatically encloses the string in single quotation marks to avoid confusing the shell. Putting the keyword `dontquote` on the next line after the keyword `command name` turns off this quoting feature for the entire dialog (this variant is not shown).

◆ **Note** Some UNIX commands insist that no spaces come between an option and its argument on a command line. In these cases, you must include control characters in the prefix definition to remove the spaces normally inserted. This is indicated in the listings by a circumflex (^) before a character. For example, `^Y` indicates CONTROL-Y, and is placed just after an option to remove the space between the option and its argument. ◆

Important Control characters do not normally print well; consequently, printouts of your dialog scripts might not show all the characters that are actually there. A circumflex followed by a letter is *not* a substitute for a control character.

Listing 13-8 shows various ways that text input is translated to the command line. The code in the top bold area of the listing formats input text on the command line with a space between the input text and the prefix. The code in the plain area between the bold areas defines a control having no space between the prefix and the text because of the CONTROL-Y at the end of the prefix. You also can remove the space before a prefix by using a CONTROL-H before the letter of the option. Each of these control characters can be used only once per option, though both can be used on a single option. The code in the bottom bold area of the listing shows how you can put several arguments having the same prefix on the command line, using the keyword `stringlist` rather than `string`.

Listing 13-8 Text box example script

```
command name "gamefinder"
help "This example demonstrates text boxes."
row {
    option name "Before you play, the GameMaster needs to know:"
        text                                    /* first control */
    }
row {}                                  /* dummy row for spacing */
row {
    option name "Your name:"            /* second control */
        prefix "-N"
```

```
    help "[-N] This enters your name."
    string
option name "Your age:"                    /* third control */
    prefix "-#^Y"     /* use Control-Y for spacing */
    help "[-#]  Your age determines the play level."
    string
option name "Games desired:"         /* fourth control */
    prefix "-T^Y"     /* use Control-Y for spacing */
    help "[-T] Specify all the games you want to try."
    stringlist
}
```

Text

Listing 13-8 also shows the use of the keyword `text` on the line labeled "first control." This control does not allow input, but simply places text in the dialog. Unlike controls that allow input, you don't specify the keywords `prefix` or `help`.

Buttons

With dialog buttons the user can

- open additional windows that allow access to files on which to operate and directories in which to save files
- open a nested dialog box, allowing choices of additional options

Dialog buttons are different from radio buttons, which select between mutually exclusive actions. Because dialog buttons have a different function, they are a different shape. Figure 13-12 shows examples of both types of dialog buttons; their names indicate their purpose. If the user clicks the Save a File button, a second dialog box appears (see Figure 13-13) and shows the standard Macintosh file dialog box for selecting a new filename. After the user selects a file, the original dialog box reappears (see Figure 13-14; note the filename in the Command Line box). If the user clicks the Redirection button, the dialog shown in Figure 13-15 comes up, allowing a choice of redirection options.

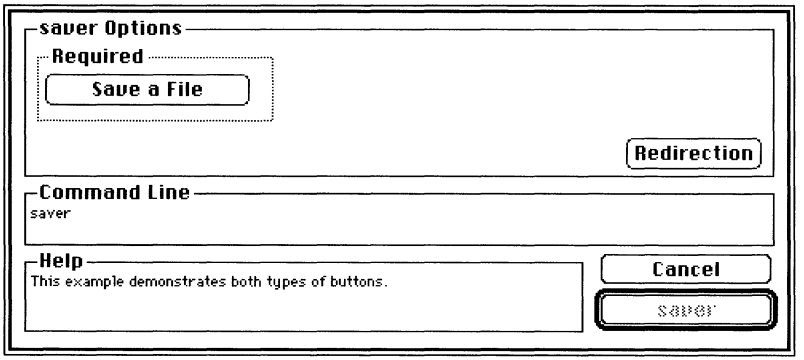

Figure 13-12 Button example: Initial dialog box

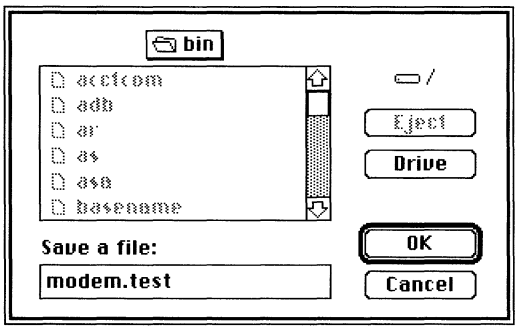

Figure 13-13 Button example: Save a File dialog box

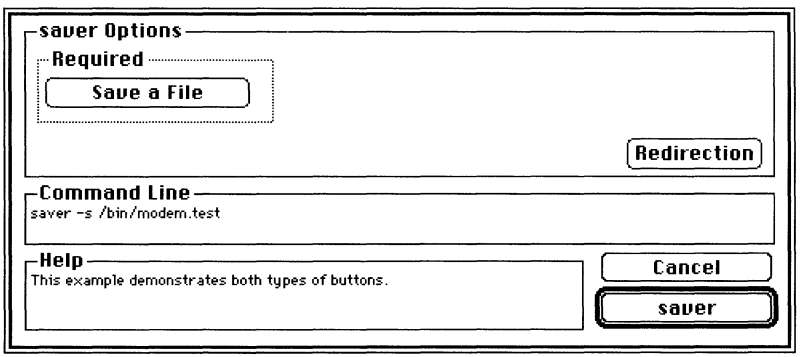

Figure 13-14 Button example: Save a File control was selected

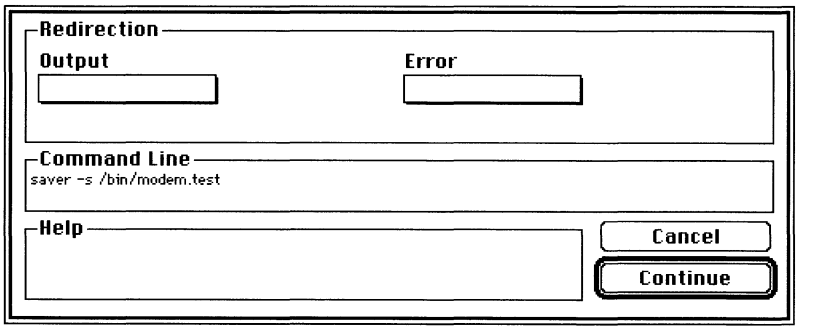

Figure 13-15 Button example: Redirection subdialog box

With dialog buttons you can call file dialogs or call a subdialog. With dialog buttons you don't have to use the keyword `prefix`, but it is good practice to always use the keyword `help` with them, though it is not required. You create file dialog buttons by putting one of the following keywords after the keyword `option name`:

```
file          dirlist
filelist      dirsandfiles
newfile       filesanddirs
directory
```

The purpose of each keyword is listed in Table 13-1. (A complete list of keywords can be found in the section "Keywords.")

Table 13-1 File dialog keywords

Keyword	Description
`file`	Presents the single file choice menu
`filelist`	Presents the file list choice menu
`newfile`	Presents the new file creation menu
`directory`	Presents the single directory choice menu
`dirlist`	Presents the directory list choice menu
`dirsandfiles`	Presents the file/directory choice menu; same as `filesanddirs`
`filesanddirs`	Presents the file/directory choice menu; same as `dirsandfiles`
`outpopup`	Presents the standard output pop-up menu
`errpopup`	Presents the standard error pop-up menu

To redirect either the standard or error output, use the keywords `outpopup` and `errpopup`. You can use `outpopup` alone; however, to use `errpopup`, you must also use `outpopup`.

To create dialog buttons that open a subdialog box, use the keyword `dialog name`. You must place this keyword after the close of a row definition (as is shown in the lower bold area of Listing 13-9). Define the name of the button with a text string (between double quotation marks) following the keyword. You can put a maximum of six buttons in a column.

Listing 13-9 shows the script that produced the dialogs in Figures 13-12 through 13-15. The first button in the script (in the upper bold area of the listing) calls a file dialog, in this case to create a new file. The first button control is followed by two dummy columns to ensure that the button does not extend the entire width of the dialog. The second button (in the lower bold area of the listing) opens a subdialog whose only components are the redirection pop-up menus.

These figures also illustrate the effects of a new keyword, `required`, found within the first control. The keyword `required` has the effect of disabling the command button until a file is selected (note the difference in the appearance of the button named *saver* between Figures 13-12 and 13-14). The keyword `required` can be used only in the first dialog of a script. The keyword `name` is used to place a box around the required control to notify the user to complete this control (see Listing 13-7). These keywords are discussed further in the next section, "Dependencies."

Listing 13-9 Button example script

```
command name "saver"

help  "This example demonstrates both types of buttons."
row {
    column {

        name "Required"          /* let user know about this */
        option name "Save a file:"          /* first button */
            prefix "-s"
            help "[-s] Saves to chosen name."
            newfile                  /* get a new file */
            required            /* HAVE to get a new file */
```

```
        }
    column {}                   /* dummy column for spacing */
    column {}                   /* dummy column for spacing */
    }
dialog name "Redirection"                   /* second button */
row {
    option name "output"
        outpopup
    option name "error"
        errpopup
    }
```

Dependencies

Controls can be selectively enabled, depending on the selection state of some other control. Controls that are disabled appear in gray (and are said to be *dimmed*); once the enabling dependency is satisfied, the control appears in black. Users can also be required to select an option.

Figures 13-16 and 13-17 show a control dependency example. In Figure 13-16, the "Pick a card" control is selected (it is the default) so the "Card name" control is enabled, while the "Suit" controls are disabled. In Figure 13-17, this order is reversed; the "Suit" control is now enabled, while the "Card name" control is disabled.

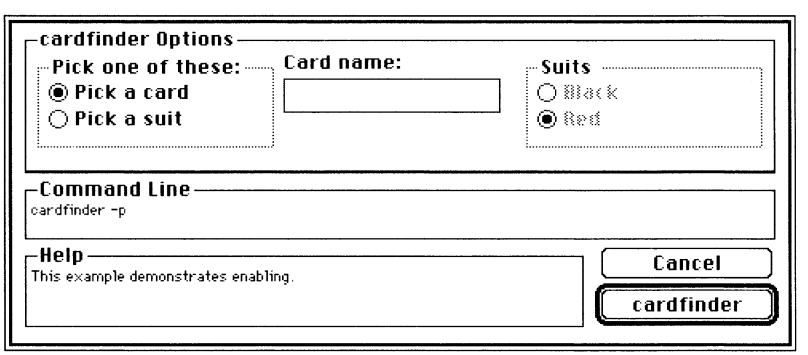

Figure 13-16 Dependencies example: First control selected

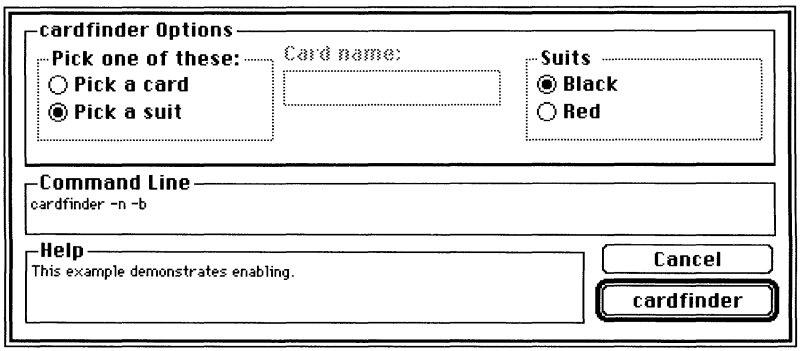

Figure 13-17 Dependencies example: Second control selected

Controls without dependencies are enabled by default; controls with dependencies are disabled by default. To disable a control, simply make its enabling dependent on another control by using the keyword `enables`. You can enable controls in two ways:

- Specify the prefix that must be in effect (showing in the Command Line box). The prefix must be identical to that used in the control specification, including any control characters within the prefix.
- Specify the control option name (the quoted text following the keywords `option name` or `name`). For example, radio buttons are enabled as a set by enclosing them in a named box and putting the name of the box in double quotation marks after the keyword `enables`.

For enabling dependencies to work, all dependent controls must be in the same dialog box. If necessary, place dependent controls together in a subdialog and enable a button that allows the user access to that subdialog.

You can require that users select an option by using the keyword `required`. The keyword `required` can be used only on the first dialog of a script. It is helpful to the user to enclose any required controls in a box named *Required* (see the examples in Figures 13-12 and 13-14).

Listing 13-10 shows the script used to created the dialogs shown in Figures 13-16 and 13-17. The first control enables the third control (see the first and third bold areas of the listing), while the second enables the grouped fourth and fifth controls (see the second and fourth bold areas of the listing). The first control enables a single control by prefix; you can use this method for all kinds of controls. The second control enables the grouped controls by name. Use this method for specifying a set of radio buttons, for individual buttons, and for controls with blank prefixes.

Listing 13-10 Dependencies example script

```
command name "cardfinder"

help "This example demonstrates enabling."
row {
    column {
    name "Pick one of these:"
    radio buttons {
        option name "Pick a card"                /* first control */
            prefix "-p"
            help "[-p] This allows selection of a card."

            enables "-c^Y"        /* enable a single control */

        option name "Pick a suit"               /* second control */
            prefix "-n"
            help "[-n] Select a suit."

            enables "Suits"    /* enable a group of controls */

                  }
           }
    option name "Card name:"                     /* third control */

        prefix "-c^Y"

        help "[-c] This enters the card name."
        string
    radio buttons {                /* set up a group of controls */

    name "Suits"

        option name "Black:"                    /* fourth control */
            prefix "-b"
            help "[-b]  Select from black suits."
        option name "Red:"                       /* fifth control */
            prefix "-r"
            help "[-r]  Select from red suits."
                  }
    }
```

The order that options appear on the command line can be specified, in reverse order, by using the keywords `last1`, `last2`, and so on. An option with the keyword `last1` appears last on the command line. An option with the keyword `last2` appears next to last, and so on. This feature can be used within a dialog box, and is nested across dialog boxes. For example, the `last1` specification of an option in the first dialog box is put on the command line after an option with the `last1` specification in any subdialog boxes.

Boxes

Outline boxes can be defined by using the keywords `box` or `name`. Each draws a box around a control or group of controls; the keyword `name` inserts a name at the top left of the box to identify its contents. The name can be as long as you like. However, if it is longer than the box, it overwrites the next column. Named boxes can be used to enable a group of radio buttons (see Figures 13-16 and 13-17, and Listing 13-10). The width of the boxes is the same as that of the current column. You can often make a dialog look better by inserting blank columns to reduce the width of the boxes (see Figures 13-7 and 13-10).

Leniencies

Commando is fairly forgiving when it comes to specifying column definitions. It is good about automatically creating columns, and usually the first column specification in a multiple column set does not need to be explicit. Radio buttons are automatically put into their own column. Commando is also reasonably well behaved as long as you don't try to put more than seven controls in a column (explicit or implicit).

Keywords

Table 13-2 alphabetically lists the keywords used in Commando.

Table 13-2 Commando keyword reference

Keyword	Description
box	Puts an outline box around a control or group of controls.
column { }	Contains the contents of a column.
command name "*name*"	Sets the name of the command in the invocation button.
dialog name "*name*"	Sets the name for a nested dialog box and the button to access it.
directory	Presents the single directory choice menu.
dirlist	Presents the directory list choice menu.
dirsandfiles	Presents the file/directory choice menu. Same as filesanddirs.
disabled	Obsolete keyword.
dontquote	Turns off the quoting mechanism for text input. Affects all text fields in a dialog script.
enables "*specifier*"	Enables other controls to be used. The control to be enabled is specified by its prefix.
errpopup	Presents the standard error redirection menu.
file	Presents the single file choice menu.
filelist	Presents the file list choice menu.
filesanddirs	Presents the file/directory choice menu. Same as dirsandfiles.
help "*help string*"	Sets the help message for this section.
last1..n	Used to specify the order of controls. last1 indicates the last option on the command line. last2 is the next-to-last, and so on.
name	Puts a named outline box around a control or group of controls.
newfile	Presents the new file creation menu.
number	Obsolete keyword.
option name "*name*"	Sets the name of checkboxes and/or buttons. Required for each control.
outpopup	Presents the standard output redirection menu. Required if errpopup is used.
prefix "*prefix string*"	Adds *prefix string* to the command line.
radio buttons { }	Defines a set of radio buttons. The braces enclose the set of controls.
required	One of the controls referenced by this keyword must be selected.
row { }	Contains the contents of a row.
string	Allows string input. The input box string width is the width of the current column.
stringlist	Allows several string inputs. The input box string width is the width of the current column.
text	Displays the control name as text.

Creating Commando dialogs

Creating new Commando dialogs is a three-step process. First, you write a new script. This usually involves copying a script that has controls similar to the ones you want to use, then modifying it to fit your application. Second, you test and, if necessary, debug the script. Third, you make the script read-only and move it to one or more places so it can be invoked by all the users on the system. If necessary, you can compile the script into a resource.

As an introduction to the process of creating dialogs, the following section examines how dialogs are invoked.

Invoking Commando dialogs

To invoke Commando from CommandShell, you can use two methods. Enter

`cmdo` *commandname*

on the command line, or type

commandname

on the command line and choose Commando from the Edit menu. The keyboard shortcut for this second method is

commandname COMMAND-K

When Commando starts, it first searches the path listed in the variable `$CMDODIR` for resources, then for dialog scripts. After that, Commando searches for resources, then dialog scripts, in the directory in `/mac/lib/cmdo` that have the same first letter as the command name you are invoking. (For example, if the invoked command is `lpr`, Commando searches the directory `/mac/lib/cmdo/l`.) Finally, Commando searches your `$PATH` variable for resources (this might result in a long search if `$PATH` includes many directories).

Make sure that the commands on the command line created by your dialog script are locatable by the shell. The normal command search path is contained in the `$PATH` shell variable. By default, this variable is set to `/bin:/usr/bin:/usr/ucb:/mac/bin:`, though this might be changed by system initialization files (such as `.profile` or `.login`).

Commando is also invoked when you double-click a UNIX application, utility, or shell script icon. This method is not efficient when you are testing dialog scripts.

Writing Commando dialogs

Although the Commando script language is reasonably straightforward, it is not foolproof. The Commando scripts that reside on each A/UX system (in `/mac/lib/cmdo/*/*`) have all been debugged and tested. Consequently, you can save time if you simply modify a script that already exists instead of trying to write your own script from scratch. This is especially true because some scripts use nonprinting control characters to enable controls, and such scripts are sometimes difficult to debug from printouts.

Testing Commando dialogs

Commando dialogs are easy to test, even when the script file is still open. When Commando is searching for script files, it searches the directories listed in the section "Invoking Commando Dialogs." Therefore, once you have written your script, simply place it in the directory within `/mac/lib/cmdo` that has the same first letter as the name of your script. The file should have read permission for your users. If you've copied and modified a file that already existed, you probably don't need to change the permissions. To set the permissions so that the file is readable by everyone, use the command line

`chmod 444` *scriptname*

If you are using TextEditor to edit a Commando file, simply save the file (you don't have to close it) in the appropriate directory within `/mac/lib/cmdo`. Commando interprets and runs the last saved version of your script. If it doesn't perform or look quite right, simply edit the file, save it again, and reinvoke the script using one of the command lines discussed earlier in this chapter.

Compiling Commando dialogs

Compiling a script into a resource file allows you to customize its appearance. Various attributes, such as the size of dialog boxes and the shape of controls, can be modified using the Commando resource editor available in MPW.

To create a Commando resource, use the command line

`cmdo` *scriptname* `-r -n -o` *outputfile*

This creates a resource file with the name *outputfile*. Move the file into a directory common to all users' `$PATH` variable, such as `/usr/bin`, so that all users can access it. After the file is moved, it must be renamed to *scriptname* so that Commando can locate the source dialog.

The command, the script, and the compiled resource must all have the same name (the resource file has a leading `%`).

Dialog design guidelines

This section offers general guidelines to assist you in planning your Commando dialogs. It is not meant to be authoritative, but does present what has been found to work best. If the needs of your applications demand it, you are free to do anything you want; but keep in mind that one of the things that makes the Macintosh so easy to use is its consistency of interface. Your design goal is to help users find choices where they expect to find them, instead of having to hunt for them. You can find many helpful hints in *Human Interface Guidelines: The Apple Desktop Interface*.

Dialog layout guidelines

Generally, it should be easy for the user to see what information is required before a command can be run and what controls are currently selected.

When a script calls for nested dialog boxes, all required arguments, as well as the most frequent or useful arguments, should be in the first dialog box. In general, try to reduce the number of dialogs to a minimum. Ideally, the user should be able to see everything in one dialog, so that it is immediately clear from the dialog box which controls have been chosen.

The layout of controls within a dialog should correspond to the direction people normally read. Required arguments, if any, should be distinguished from optional arguments and presented in the first part of the first dialog page. The most important or frequently used arguments should follow after the required arguments. For example, in France people usually read left to right and top to bottom, so the layout of the dialog and controls should follow this pattern.

Use boxes to group similar items. Boxes can separate columns, portions of columns, or clusters of buttons. Boxes do not have to be labeled, though labels are often useful.

Buttons to select files or directories (or both) should be placed on the first dialog page when possible. Use the keywords `last1`, `last2`, and so on to permit this arrangement.

Normally, each dialog item corresponds to a single control or argument. In some cases, however, a command can have one or more commonly used group of controls. In these cases, some of the dialog items might correspond to control clusters. Note that the user should still be able to select all controls individually.

Use the keyword `string` if the possible values are infinite. If the number of values is a small, finite number, try to use radio buttons.

There are several standards for subdialog names:

- Subdialogs containing only Output and Error pop-up menus should be labeled "Output & Error."
- If a dialog contains only one subdialog of unrelated options, that subdialog should be labeled "More options." If the options are closely related, that relationship can be used to name the subdialog.
- If a dialog contains several subdialogs containing unrelated options, these subdialogs should be named "Options 1," "Options 2," and so on.

Dialog aesthetics

Try to avoid mixing control types (checkboxes, radio buttons, text boxes, and buttons). Try to make the dialog page look balanced. With few exceptions, dialog boxes look best with two columns per row. Use empty columns for spacing to prevent a column from appearing too wide.

Don't juxtapose unrelated sets of radio buttons. Remember that the first radio button in a cluster is turned on by default. Take care to choose a default that is reasonable. It is often a good idea to add a button to a cluster of controls to represent the default action.

Descriptive information

The labels associated with a dialog item should be understandable by the UNIX-naive user whenever possible. Options should be described in terms of the results that the user will see, rather than in terms of the underlying UNIX concepts.

Filename arguments should be specified by their function or role. For example, use "Files to be searched" rather than "Input."

Always try to show the effects of defaults. One example is to label the pop-up menus for output files "Output to" so that the default behavior is displayed on the screen.

Put useful information on the screen if it doesn't lead to clutter. For example, the UNIX command `date` takes as an argument a string formatted `mmddhhmm[yy]`. This format is small, useful, and easy to forget. It can be placed just above the text box where the user can refer to the format when entering the date. Examples of more extended information should be placed in the help message.

The help messages should expand on the text in the upper portion of the dialog box to provide information and, where possible, examples. Don't simply repeat the control text for the help message. If you can't think of anything else to add, rephrase the control text in case the user didn't understand the original text. When the user has to type in something, give examples of common usages.

Index

%nonassoc keyword 3-24
%prec keyword 3-25
%right keyword 3-24
%type keyword 3-37, 3-38
.DEFAULT target 7-10
.IGNORE target 7-10
.MAKESTOP target 7-10
.PRECIOUS target 7-10
.SILENT target 7-10
/usr/lib/MakeRules file 7-9

A

accept action 3-14
addch function 12-8
addch routine 12-4, 12-28
addstr routine 12-16, 12-28
admin command 8-3, 8-22
alphabetic keyword reference 13-29
ar command 6-3
archive 6-3
attroff routine 12-32
attron routine 12-32
attrset routine 12-8, 12-32
auto statement 11-10
awk 9-1 to 9-61
  action 9-5
  action block 9-4
  actions 9-20 to 9-34
arrays 9-38
  special for loop 9-40
assignment operators 9-37
BEGIN pattern 9-5, 9-19
braces 9-54
built-in
  functions 9-47
  numerics 9-49
  variables 9-11, 9-40
  comments 9-5
conditions 9-22
  expressions 9-22
data structures 9-35 to 9-41
data type determination 9-46
directing output 9-34
END pattern 9-5, 9-19
expressions 9-21, 9-41 to 9-50, 9-60
  assignment 9-61
field separator 9-6
flow of control 9-23
functions 9-22, 9-57
identifiers 9-52
input processing 9-11
invocation 9-7
lexical conventions 9-50
looping constructs 9-23
matching operation 9-15
numeric constants 9-50 to 9-51, 9-55
operation 9-3
operators
  arithmetic 9-42
  assignment 9-42
  logical 9-44
  pattern-matching 9-44
  relational 9-43
  symbols 9-22
options 9-6
pattern 9-5
pattern-seeking operation 9-15
patterns 9-14 to 9-20
  expressions 9-15
predefined variables 9-51
primary expressions 9-55
printing
  output 9-30, 9-31 to 9-34
  variables 9-32
program components 9-21
reading input 9-29
records 9-54
regular expression 9-17
report generation 9-27
reserved function names 9-51
reserved keywords 9-51
separators 9-53, 9-54
  field 9-53
  record 9-53
shell interactions 9-9
shell scripts 9-9

awk *(continued)*
special characters 9-17
string constants 9-51, 9-56
system command 9-34
terms 9-58 to 9-60
binary 9-58
unary 9-59
tokens 9-50
variable initialization 9-37
variables 9-35, 9-56 to 9-57
incremented 9-59

B

backslash escape 5-9, 5-10
backspace escape 5-9, 5-10
baudrate function 12-33, 12-47
bc program 11-1 to 11-20
arrays 11-9
assignment statements 11-12
automatic variables 11-8
comments 11-5
constants 11-6
control statements 11-13
defining functions 11-7
exiting 11-4
expressions 11-15 to 11-17
function calls 11-7
global variables 11-9
I/O base 11-17 to 11-19
identifiers 11-6
keywords 11-6
operators 11-16
program files 11-4
program syntax 11-5
relational operators 11-14
scale 11-19 to 11-20
statements 11-10 to 11-11
syntax 11-3
usage 11-3
bdiff command 6-2
beep routine 12-17, 12-33
BEGIN pattern 9-5
box keyword 13-28, 13-29
box routine 12-30
boxes 13-19, 13-28
break statement 9-23, 11-10

C

C flowgraph, cflow 2-2
C-language, escapes 5-9
C preprocessor, cpp 2-3
calling dialogs 98
cb 2-2
CBREAK mode 12-23
cbreak routine 12-3, 12-5, 12-23
cdc command 8-26
cflow command 2-2
changequote macro 4-4
checkbox 13-16
clear routine 12-28
clearok routine 12-16, 12-20
close function 9-30
clrtobot routine 12-28
clrtoeol routine 12-28
cmdo command 13-30
COFF symbol table 2-5
column keyword 13-7, 13-29
comb command 8-26
comm command 6-2
command name 13-6, 13-9
command name keyword 13-29
Commando
dialog boxes 13-4
keyword reference alphabetic 13-29
script language 13-5
comments 13-7
awk 9-5
bc 11-5
make 7-12
yacc 3-6
compiling dialogs 13-32
condition 9-22
continue statement 9-23
control
characters 13-20
dependencies 13-25
examples 13-16
cpp, with make 7-51
cpp language 2-3
creating Commando dialogs 13-30
ctags command 2-4
curses 12-1 to 12-67
additional terminals 12-40
CBREAK mode 12-23
constants 12-18
curses.h file 12-19
delays 12-34
examples 12-47 to 12-67
function keys 12-6
highlighting 12-8
initialization 12-19
input 12-5
terminal 12-31
terminating 12-17
lower-level functions 12-35
mini-curses 12-43
multiple terminals 12-11, 12-40
operation 12-39 to 12-47
options 12-20
output 12-4, 12-25
portability 12-46
functions 12-33
routines 12-18
scrolling 12-43
special keys 12-42
structure 12-18
terminal mode 12-23
terminfo 12-35
terminfo usage 12-13
TTY-mode functions 12-45
typeahead check 12-45
usage 12-3
variables 12-18
video attributes 12-32, 12-41
windows
attributes 12-8 to 12-9
manipulation 12-24
multiple 12-10
writing 12-27
curses.h file 12-19

D

dc program 10-1 to 10-10
  arrays 10-7
  base numbering 10-4
  commands 10-4
  input conversion 10-4
  operators 10-3
  output commands 10-5
  programming 10-9 to 10-10
  reference 10-8
  registers 10-7
  registers, internal 10-6
  scale 10-5
  scale rules 10-5
  stack commands 10-6
  subroutine definitions 10-6
  syntax 10-2
  usage 10-2
decr function 4-9
define function 4-3, 11-10
delay_output routine 12-39
delch function 12-16
delch routine 12-29
deleteln function 12-16, 12-29
delta command 8-5, 8-28
delwin routine 12-25
dependency 7-4
description file (see makefile) 7-13
dialog boxes 13-4
  aesthetics 13-33
  design 13-32
  invoking 13-30
  layout 13-5, 13-32
  text in 13-33
dialog name keyword 13-23, 13-29
diff command 6-2
diff3 command 6-2
diffdir command 6-2
diffmk command 6-2
directory keyword 13-23, 13-29
dirlist keyword 13-23, 13-29
dirsandfiles keyword 13-23, 13-29
disabled keyword 13-29
disambiguating rule 3-20
divert function 4-11
divnum function 4-13
dnl macro 4-14
dontquote keyword 13-19, 13-29
doupdate routine 12-26
draino routine 12-34
dummy column 13-12
dumpdef macro 4-16

E

ECHO function 5-20
echo routine 12-23
enables 13-26
  keyword 13-29
enabling
  by name 13-27
  by prefix 13-27
END pattern 9-5
end-marker token 3-8, 3-13
endwin routine 12-4, 12-19
erase routine 12-28
erasechar function 12-33, 12-46
error
  action 3-14, 3-16
  redirection 13-23
  token 3-13, 3-27
errpopup keyword 13-23, 13-29
errprint macro 4-16
eval function 4-9
exit statement 9-24
exp function 9-58

F

filelist keyword 13-23, 13-29
files
  COFF sections 6-2
  comparing 6-2
  dialog keywords 13-23
  finding 6-2
  keyword 13-23, 13-29
  manipulation tools 6-1
  version 6-3
filesanddirs keyword 13-23, 13-29
find command 6-2
flash routine 12-17
floating-point constants, in lexical analysis 3-47
flushinp routine 12-34, 12-47
for loop 9-25, 9-40
for statement 11-13
  bc 11-12
function, finding definition 2-4

G

get command 8-6, 8-30
getch routine 12-6, 12-31
getline command 9-24, 9-30
getstr routine 12-6, 12-32, 12-46
getyx routine 12-30
goto action 3-15
grammar rules 3-3
  left recursive 3-31
  right recursive 3-32
gsub function 9-49

H

help 13-6, 13-9
help command 8-6, 8-39
help keyword 13-29
help messages 13-34
  length 13-9

I, J

ibase function 11-17
idlok routine 12-3, 12-20
if statement 9-24
  bc 11-12
ifdef macro 4-6
ifelse macro 4-8
inch routine 12-30
include function 4-10
incr function 4-9
index function 9-52
index macro 4-15
initscr routine 12-3, 12-19

input routine 5-23
insch function 12-16, 12-29
insertln function 12-16, 12-29
int function 9-57
intrflush routine 12-22
invoking dialogs 13-30
iodlk routine 12-39

K

keypad routine 12-6, 12-21, 12-42
keyword
  box 13-18, 13-28, 13-29
  column 13-7
  command name 13-6, 13-9
  dialog name 13-23
  directory 13-23
  dirlist 13-23
  dirsandfiles 13-23
  dontquote 13-19
  enables 13-26
  errpopup 13-23
  file 13-23
  filelist 13-23
  filesanddirs 13-23
  help 13-6, 13-9
  last1 13-28
  name 13-18, 13-28, 13-29
  newfile 13-23
  outpopup 13-23
  prefix 13-10
  required 13-24, 13-26, 13-29
  row 13-7
  string 13-20
  stringlist 13-20
  text 13-21
killchar routine 12-34, 12-46

L

last1 13-28
  keyword 13-29
leaveok routine 12-21
left association 3-19
len macro 4-14
length function 9-48, 11-10
lex 5-1 to 5-31
  actions 5-9 to 5-26
  alternation 5-14
  ambiguous rules 5-18 to 5-19
  arbitrary characters 5-9
  character classes 5-7
  character set 5-7
  compilation 5-27
  context sensitivity 5-14
  definition expansion 5-12
  definitions 5-10
  examples 5-27 to 5-29
  expressions
    operators 5-31
    optional 5-13
    regular 5-12
    repeated 5-13
  flags 5-16
  I/O routines 5-23
  null statement 5-20
  operators 5-9
  repetition character 5-20
  repetitions 5-12
  rules 5-12
  start conditions 5-17
  substitution strings 5-11
  summary 5-29
  syntax 5-6
  usage 5-3
  variables 5-10
  yacc usage 5-4
library 6-3
log function 9-57
longname routine 12-20, 12-44
look-ahead token 3-14

M

m4 macro processor 4-1 to 4-20
  arguments 4-7
  arithmetic
    functions 4-9
    operators 4-9
  I/O 4-10
  invocation 4-3
  macro
    definition 4-3
    summary 4-19
  printing 4-16
  quoting 4-5
  recursive definitions 4-17
  string manipulation 4-14
  system commands 4-16
Macintosh dialog boxes 13-3
macro
  arguments 4-7
  definition 4-3
  replacement 4-7
make
  archive libraries 7-40 to 7-42
  attributes 7-38 to 7-39
  built-in macros 7-31
  built-in targets 7-10
  colons 7-11
  combining commands 7-19
  command syntax 7-5
  commands 7-12
  comments 7-13
  default commands 7-19
  default rules 7-28
  dependency 7-4
  dependency statement 7-11
  description file 7-8. *See also* makefile
  dynamic dependency parameters 7-16
  Dynamic Include File Dependency Generation (DIFDG) 7-50
  environment variables 7-28
  errors 7-18
  include directives 7-51
  include lines 7-13
  internal macros 7-15
  macros 7-29
    definitions 7-13
    expansion 7-30
    expansion characters 7-33 to 7-34
    setting defaults 7-26
    testing 7-37
    translation 7-48

makefile 7-3, 7-20
  entries 7-9
  mnemonic targets 7-47
  operation 7-28
  options 7-6, 7-18
  precedence 7-35
  predecessor trees 7-45
  printing command names 7-18
  rules 7-21
    user defined 7-9
  SCCS files 7-42 to 7-43
  special characters 7-33 to 7-34
  suffix list default 7-27
  suffixes 7-20
  target 7-4
  transformation rules 7-21 to 7-27
  usage 7-3
  walking directory trees 7-44
make program 7-8 to 7-52
MAKEBDIR macro 7-31
MAKECDIR macro 7-31
makefile 7-3, 7-20
  writing 7-3
MAKEFLAGS environment variable 7-35
MAKEFLAGS macro 7-31
MAKEGOALS macro 7-31
MAKELEVEL macro 7-31
match function 9-49
meta routine 12-21
mini-curses 12-43
move routine 12-3, 12-27
mvcur routine 12-35
mvinch function 12-16
mvwin routine 12-25

N

name 13-18
name keyword 13-28, 13-29
napms routine 12-34
newfile keyword 13-23, 13-29
newline escape 5-10
newpad routine 12-24
newterm routine 12-12, 12-20
newwin routine 12-10, 12-24
next statement 9-24
nl routine 12-23
nm command 2-5
node function 3-11
nodelay routine 12-5, 12-21, 12-46
noecho routine 12-5
nonterminal symbol 3-3, 3-6, 3-7
null statement 5-20
number keyword 13-29

O

obase function 11-17, 11-18
octal dump 2-5
od command 2-5
option
  dependencies 13-25
  leniencies 13-28
  name 13-9
  name keyword 13-29
  order 13-28
  type
    checkboxes 13-16
    radio buttons 13-17
    text 13-20 to 13-21
    text box 13-19

P

outpopup keyword 13-23, 13-29
output redirection 13-23
output routine 5-23
overlay routine 12-25
parse trees 3-11
parser 3-3
prefix keyword 13-10, 13-29
prefresh routine 12-26
print command 9-30, 9-31 to 9-32
  variables 9-32
printf command 5-20, 9-31, 9-33, 9-35
printw routine 12-4, 12-30
prof command 2-2
profile data, prof 2-2
program structure, cb 2-2
prs command 8-39
putp routine 12-39

Q

quit statement 11-11
quotes, macros 4-5

R

radio buttons keyword 13-18, 13-29
raw routine 12-23
reduce action 3-14 to 3-15
refresh mode 12-25
refresh routine 12-4
REJECT action 5-25
repetition character 5-20
required keyword 13-24, 13-26, 13-29
resetty routine 12-24
return statement 11-8, 11-11
right association 3-19
rmdel command 8-41
row keyword 13-7, 13-29

S

sact command 8-42
scalars 3-46 to 3-47
scale function 11-3, 11-11
SCCS 8-1 to 8-44
  administering 8-9
  arguments 8-16
  branch deltas 8-14
  change comments 8-26
  command summary 8-22
  commands 8-16
  comments 8-24 to 8-25, 8-28 to 8-29
  delta
    combination 8-26
    numbering 8-13
    removal 8-41
  descriptive text 8-25
  diagnostics 8-17
  ERROR 8-17

SCCS *(continued)*
files 8-7 to 8-15
accounts 8-42
arguments 8-16
auditing 8-12
check characteristics 8-43
comparison 8-43
corrupt 8-13
creating 8-3, 8-22
format 8-12
temporary 8-17, 8-27
flags 8-17, 8-23
group projects 8-9
ID keywords 8-20
identification string 8-3
keyletters 8-16, 8-37
keywords 8-29
modification request (MR) numbers 8-24, 8-28
new versions 8-27
on-line explanations 8-39
on-line information 8-6
printing 8-39
protection 8-7
restoring version 8-38
retrieving versions 8-5, 8-30, 8-31
SID determination 8-35
`sccsdiff` command 8-43
screen I/O, see curses 12-1
script structure 13-6, 13-8
`scroll` routine 12-30
`scrollok` routine 12-22
`set_term` routine 12-20
`setscrreg` routine 12-22, 12-43
`setuid` bit 8-11
`setupterm` routine 12-14, 12-36
`SHELL` variable 7-29
`shift` action 3-14
shift/reduce conflict 3-20
SID. *See* SCCS, identification string
`sinclude` function 4-10
`size` command 6-2
`sort` command 9-28
`split` function 9-47
`sprintf` function 9-49
`sqrt` function 9-57, 11-10
`standend` routine 12-32
`standout` routine 12-32
`stdscr` routine 12-10
`string` keyword 13-19, 13-29
`stringlist` keyword 13-19, 13-29
strings
quoting 4-5
sharing 2-4
`substr` macro 4-15
`substring` function 9-48
`subwin` routine 12-24
`syscmd` macro 4-16
`system` command 9-34

T

tab escape 5-10
target 7-4
`termcap` database 12-2
terminal symbol 3-3
`terminfo`
database 12-2, 12-35
usage 12-13
terms 9-58
testing dialogs 13-31
`text box` keyword 13-19 to 13-20
`text` keyword 13-21, 13-29
tokens 3-3, 9-50
error 3-27
look-ahead 3-14
names 3-13
number 3-12, 3-13
`touch` command 7-47
`touchwin` routine 12-25
`tparm` routine 12-38
`tputs` function 12-15, 12-38
`typeahead` routine 12-22

U

`undefine` macro 4-6
`undivert` function 4-11
`unget` command 8-38
`unput` routine 5-23

V

`val` command 8-43
`version` command 6-3
`vidattr` routine 12-38
`viputs` routine 12-38
`VPATH` macro 7-23

W

`what` command 8-43
`while` loop 9-24
`bc` 11-12
`window` structure 12-4
`wrefresh` routine 12-10
writing dialogs 13-31

X

`xstr` command 2-4

Y, Z

`y.output` file 3-16, 3-26
`yacc` 3-1 to 3-56
actions 3-8 to 3-11, 3-34
ambiguity 3-19
arithemetic expressions 3-23
comments 3-6
conflict resolution 3-26
conflicts 3-19
declarations 3-11, 3-13
environment 3-29
error handling 3-4, 3-27
escapes 3-7
examples 3-38 to 3-55
floating-point constants 3-46
grammar rules 3-6
hints 3-31
input 3-31
`lex` usage 5-3
lexical
analysis 3-12 to 3-13
considerations 3-32
library 3-30
literal characters 3-4

null character 3-7
parser
  operation 3-14 to 3-18
  steps 3-14
precedence 3-23 to 3-26
  rules 3-25
recursion 3-31 to 3-32
reserved words 3-34
specifications 3-6 to 3-8
  files 3-4
  tokens 3-12, 3-13
  type checking 3-37
  union 3-36
  usage 3-3
  values 3-36
`YYACCEPT` macro 3-34
`yychar` variable 3-30
`yyclearin` statement 3-29
`yydebug` variable 3-30
`yyerror` function 3-30
`YYERROR` macro 3-34
`yyerrorok` statement 3-28
`yyleng` count 5-21
`yyless` routine 5-22
`yylex` function 3-12
`yymore` routine 5-22
`yyparse` function 3-29
`yywrap` routine 5-24

The Apple Publishing System

A/UX Programming Languages and Tools, Volume 2, was written, edited, and composed on a desktop publishing system using Apple Macintosh computers, an AppleTalk network system, Microsoft Word, and QuarkXPress. Line art was created with Adobe Illustrator. Proof pages were printed on Apple LaserWriter printers. Final pages were output directly to 70-mm film on an Electrocomp 2000 Electron Beam Recorder. PostScript, the LaserWriter page-description language, was developed by Adobe Systems Incorporated.

Text and display type are Apple's corporate font, a condensed version of ITC Garamond®. Bullets are ITC Zapf Dingbats®. Some elements, such as program listings, are set in Apple Courier, a fixed-width font.

Writer: J. Eric Akin
Developmental Editor: Scott Smith
Design Director: Lisa Mirski
Art Director: Tamara Whiteside
Production Editor: Debbie McDaniel

Special thanks to Jeannette Allen, Tom Berry, Vicki Brown, Gene Garbutt, Michael Hinkson, Kristi Fredrickson, John Morley, John Sovereign, Earl Wallace, Kathy Wallace, Kristen Webster, Laura Wirth, and Chris Wozniak.