




Programming Utilities *for the Sun Workstation®*



Credits and Acknowledgements

The chapters of this manual were originally derived from the work of many people at Bell Laboratories, the University of California at Berkeley, and other noble institutions. Their names and the titles of the original works appear here.

UNIX Programming

by Brian W. Kernighan and Dennis M. Ritchie, Bell Laboratories, Murray Hill, New Jersey.

Lint, a C Program Checker

by S. C. Johnson, Bell Laboratories, Murray Hill, New Jersey.

Make — A Program for Maintaining Computer Programs

by S. I. Feldman, Bell Laboratories, Murray Hill, New Jersey.

The M4 Macro Processor

by Brian W. Kernighan and Dennis M. Ritchie, Bell Laboratories, Murray Hill, New Jersey.

Lex — A Lexical Analyzer Generator

by M. E. Lesk and E. Schmidt, Bell Laboratories, Murray Hill, New Jersey.

Yacc — Yet Another Compiler-Compiler

by Stephen C. Johnson, Bell Laboratories, Murray Hill, New Jersey.

Source Code Control System User's Guide

by L. E. Bonanni and C. A. Salemi, Bell Laboratories, Piscataway, New Jersey.

Source Code Control System

by Eric Allman, Formerly of Project Ingres, University of California at Berkeley.

Trademarks

Multibus is a trademark of Intel Corporation.

Sun Workstation is a trademark of Sun Microsystems Incorporated.

UNIX is a trademark of Bell Laboratories.

Copyright © 1983, 1984, 1985, by Sun Microsystems.

This publication is protected by Federal Copyright Law, with all rights reserved. No part of this publication may be reproduced, stored in a retrieval system, translated, transcribed, or transmitted, in any form, or by any means manual, electric, electronic, electro-magnetic, mechanical, chemical, optical, or otherwise, without prior explicit written permission from Sun Microsystems.

Contents

Chapter 1	UNIX Programming	3
1.1.	Basics	3
1.2.	The ‘Standard Input’ and ‘Standard Output’	4
1.3.	The Standard I/O Library	6
1.4.	Low-Level Input Output	9
1.5.	Processes	14
1.6.	Signals — Interrupts and All That	19
1.7.	The Standard I/O Library	23
Chapter 2	Tools for the C Programming Language	33
2.1.	ctags — Build Index File for C Functions	33
2.2.	Lint — A C Program Checker	35
Chapter 3	Make — Maintaining Computer Programs	49
3.1.	Basic Features	50
3.2.	Description Files	52
3.3.	Using the make Command	56
3.4.	Implicit Rules	59
3.5.	Example	61
3.6.	Suggestions and Warnings	63
3.7.	Making Archive Libraries	64
3.8.	Suffixes and Transformation Rules	65
Chapter 4	Source Code Control System	71

4.1. Learning the Lingo	74
4.2. Creating SCCS Database Files with <code>sccs create</code>	75
4.3. Retrieving Files for Compilation with <code>sccs get</code>	75
4.4. Changing Files (Creating Deltas)	76
4.5. Restoring Old Versions	79
4.6. Auditing Changes	80
4.7. Shorthand Notations	81
4.8. Using SCCS on a Project	82
4.9. Saving Yourself	82
4.10. Managing SCCS Files with <code>sccs admin</code>	83
4.11. Maintaining Different Versions (Branches)	83
4.12. Using <code>sccs</code> with <code>make</code>	85
4.13. SCCS Quick Reference	89
Chapter 5 Performance Analysis	93
5.1. <code>time</code> — Display Time Used by Program	93
5.2. <code>prof</code> — Generate Profile of Program	96
5.3. <code>gprof</code> — Generate Call Graph Profile of Program	98
5.4. <code>tcov</code> — Statement-level Analysis of Program	100
Chapter 6 m4 — A Macro Processor	107
6.1. Using the <code>m4</code> Command	108
6.2. Defining Macros	108
6.3. Quoting and Comments	109
6.4. Macros with Arguments	110
6.5. Arithmetic Built-ins	111
6.6. File Manipulation	112
6.7. Running System Commands	113
6.8. Conditionals	113
6.9. String Manipulation	114
6.10. Printing	115
6.11. Summary of Built-in <code>m4</code> Macros	115

Chapter 7 Lex — A Lexical Analyzer Generator	119
7.1. Lex Source	122
7.2. Lex Regular Expressions	123
7.3. Lex Actions	126
7.4. Ambiguous Source Rules	129
7.5. Lex Source Definitions	131
7.6. Using <code>lex</code>	132
7.7. Lex and Yacc	133
7.8. Examples	133
7.9. Left Context-Sensitivity	136
7.10. Character Set	138
7.11. Summary of Source Format	138
7.12. Caveats and Bugs	140
Chapter 8 Yacc — Yet Another Compiler-Compiler	143
8.1. Basic Specifications	145
8.2. Actions	147
8.3. Lexical Analysis	149
8.4. How the Parser Works	151
8.5. Ambiguity and Conflicts	155
8.6. Precedence	159
8.7. Error Handling	162
8.8. The Yacc Environment	163
8.9. Hints for Preparing Specifications	164
8.10. Advanced Topics	166
8.11. A Simple Example	170
8.12. Yacc Input Syntax	173
8.13. An Advanced Example	176
8.14. Old Features Supported but not Encouraged	182
Appendix A SCCS Low-Level Commands	185
A.1. Low Level SCCS For Beginners	185
A.2. SCCS File Numbering Conventions	186

A.3. Summary of SCCS Commands	188
A.4. SCCS Command Conventions	189
A.5. <code>admin</code> — Create and Administer SCCS Files	190
A.6. <code>cdc</code> — Change Delta Commentary	197
A.7. <code>comb</code> — Combine SCCS Deltas	198
A.8. <code>delta</code> — Make a Delta	200
A.9. <code>get</code> — Get Version of SCCS File	204
A.10. <code>help</code> — Ask for SCCS Help	216
A.11. <code>prs</code> — Print SCCS File	217
A.12. <code>rmDEL</code> — Remove Delta from SCCS File	221
A.13. <code>sact</code> — Display SCCS Editing Activity	223
A.14. <code>sccsdiff</code> — Display Differences in SCCS Versions	223
A.15. <code>unget</code> — Undo a Previous SCCS <code>get</code>	224
A.16. <code>val</code> — Validate SCCS File	224
A.17. SCCS Files	226
Appendix B Bibliography and Credits	231

Tables

Table 3-1 Summary of Assigning Macros and Variables	57
Table 3-2 Default Suffix List for Make	59
Table 3-3 Built In Compiler Names and Options	61
Table 5-1 Control Key Letters for the time Command	95
Table 5-2 Default Timing Summary Chart	95
Table 6-1 Operators to the eval built in in m4	112
Table 6-2 Summary of Built-in m4 Macros	115
Table 7-1 Changing Internal Array Sizes in lex	139
Table 7-2 Regular Expression Operators in lex	140
Table A-1 Determination of SCCS Identification String	207
Table A-2 Identification Keywords	209
Table A-3 SCCS Files Data Keywords	219
Table A-4 Codes Returned from val Command	225

Figures

Figure 3-1	Single Suffix Rules	65
Figure 3-2	Double Suffix Rules	66
Figure 4-1	Basic SCCS Operations	73
Figure 7-1	An overview of Lex	120
Figure 7-2	Lex with Yacc	121
Figure 7-3	Sample character table.	138
Figure A-1	Evolution of an SCCS File	186
Figure A-2	Tree Structure with Branch Deltas	187
Figure A-3	Extending the Branching Concept	188

Preface

Welcome to *Programming Tools for the Sun Workstation*. This manual is a comprehensive description of the software utilities available to assist programmers generating software.

Summary of Contents

Chapter 1 — UNIX Programming describes the basics of using the UNIX† library routines and system calls.

Chapter 2 — Tools for the C Programming Language describes some of the tools available to assist C language programming.

Chapter 3 — Make — Maintaining Computer Programs describes a tool to assist in building, regenerating, and keeping up to date programs constructed from many source modules with dependencies between the pieces.

Chapter 4 — Source Code Control System describes the facilities available to manage and keep history of source code and documentation. This chapter describes the ‘high-level’ SCCS interface. There is also a ‘low-level’ SCCS interface described in appendix A — SCCS Low-Level Commands.

Chapter 5 — Performance Analysis covers tools available for determining how much resources a program consumes and how to focus in on where a program is spending its time.

Chapter 6 — m4 — A Macro Processor describes a simple macro processor that can be used as a front end to any other language processor.

Chapters 7 and 8 cover Lex — A Lexical Analyzer Generator and Yacc — Yet Another Compiler-Compiler. These two tools are valuable for constructing lexical and syntactic analyzers.

Appendix A — SCCS Low-Level Commands describes the SCCS low-level command interface and contain a summary of SCCS commands.

Appendix B — *Bibliography and Credits* — contains the bibliography, credits, and acknowledgements for the rest of this manual.

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

1

UNIX Programming

UNIX Programming	3
1.1. Basics	3
Program Arguments	3
1.2. The ‘Standard Input’ and ‘Standard Output’	4
1.3. The Standard I/O Library	6
Accessing Files	6
Error Handling — Stderr and Exit	9
Miscellaneous I/O Functions	9
1.4. Low-Level Input Output	9
File Descriptors	9
read and write	10
Open, Creat, Close, Unlink	12
Random Access — Seek and Lseek	13
Error Processing	14
1.5. Processes	14
The ‘System’ Function	14
Low-Level Process Creation — Execl and Execv	15
Control of Processes — Fork and Wait	16
Pipes	17
1.6. Signals — Interrupts and All That	19
1.7. The Standard I/O Library	23
General Usage	23
Standard I/O Library Calls	24

Character Type Checking	30
Character Type Conversion	30

UNIX Programming

This chapter is an introduction to programming on the UNIX[†] system. The emphasis is on how to write programs that interface to the operating system, either directly or through the standard I/O library. The topics discussed include

- handling command arguments
- rudimentary I/O; the standard input and output
- the standard I/O library; file system access
- low-level I/O: open, read, write, close, seek
- processes: exec, fork, pipes
- signals — interrupts, etc.

Section 1.7 — *The Standard I/O Library* — describes the standard I/O library in detail.

This chapter describes how to write programs that interface with the UNIX operating system in a nontrivial way. This includes programs that use files by name, that use pipes, that invoke other commands as they run, or that attempt to catch interrupts and other signals during execution.

The document collects material which is scattered throughout several sections of the Sun Reference Manuals (*Commands Reference Manual* and *UNIX Interface Reference Manual*[1]. There is no attempt to be complete; only generally useful material is dealt with. It is assumed that you will be programming in C, so you must be able to read the language roughly up to the level of *The C Programming Language*[2]. You should also be familiar with UNIX itself.

1.1. Basics Program Arguments

When a C program is run as a command, the arguments on the command line are made available to the function `main` as an argument count `argc` and an array `argv` of pointers to character strings that contain the arguments. By convention, `argv[0]` is the command name itself, so `argc` is always greater than 0.

The following program illustrates the mechanism: it simply echoes its arguments back to the terminal — This is essentially the `echo` command.

[†] UNIX is a trademark of AT&T Bell Laboratories.

```
main(argc, argv)    /* echo arguments */
int argc;
char *argv[];
{
    int i;

    for (i = 1; i < argc; i++)
        printf("%s%c", argv[i], (i<argc-1) ? ' ' : '\n');
}
```

`argv` is a pointer to an array whose elements are pointers to arrays of characters; each is terminated by `\0`, so they can be treated as strings. The program starts by printing `argv[1]` and loops until it has printed `argv[argc-1]`.

The argument count and the arguments are parameters to `main`. If you want to keep them around so other routines can get at them, you must copy them to external variables.

1.2. The 'Standard Input' and 'Standard Output'

The simplest input mechanism is to read from the 'standard input,' which is generally the user's terminal. The function `getchar` returns the next input character each time it is called. A file may be substituted for the terminal by using the `<` convention (input redirection): if `prog` uses `getchar`, the command line

```
tutorial% prog < filename
```

makes `prog` read from the file specified by `filename` instead of the terminal. `prog` itself need know nothing about where its input is coming from. This is also true if the input comes from another program via the *pipe* mechanism:

```
tutorial% otherprog | prog
```

provides the standard input for `prog` from the standard output (see below) of `otherprog`.

`getchar` returns the value `EOF` when it encounters the end of file (or an error) on whatever you are reading. The value of `EOF` is normally defined to be `-1`, but it is unwise to take any advantage of that knowledge. As will become clear shortly, this value is automatically defined for you when you compile a program, and need not be of any concern.

Similarly, `putchar(c)` puts the character `c` on the 'standard output', which is also by default the terminal. The output can be captured on a file by using `>`: if `prog` uses `putchar`,

```
tutorial% prog > outputfile
```

writes the standard output on `outputfile` instead of the terminal. `outputfile` is created if it doesn't exist; if it already exists, its previous contents are overwritten. A pipe can be used:


```
tutorial% prog | otherprog
```

puts the standard output of `prog` into the standard input of `otherprog`.

The function `printf`, which formats output in various ways, uses the same mechanism as `putchar` does, so calls to `printf` and `putchar` may be intermixed in any order; the output will appear in the order of the calls.

Similarly, the function `scanf` provides for formatted input conversion; it will read the standard input and break it up into strings, numbers, etc., as desired. `scanf` uses the same mechanism as `getchar`, so calls to them may also be intermixed.

Many programs read only one input and write one output; for such programs I/O with `getchar`, `putchar`, `scanf`, and `printf` may be entirely adequate, and it is almost always enough to get started. This is particularly true if the UNIX pipe facility is used to connect the output of one program to the input of the next. For example, the following program strips out all ASCII control characters from its input (except for newline and tab).

```
#include <stdio.h>

main() /* ccstrip: strip non-graphic characters */
{
    int c;
    while ((c = getchar()) != EOF)
        if ((c >= ' ' && c < 0177) || c == '\t' || c == '\n')
            putchar(c);
    exit(0);
}
```

The line

```
#include <stdio.h>
```

should appear at the beginning of each source file which does I/O using the standard I/O functions described in section 3(S) of the *UNIX Interface Reference Manual* — the C compiler reads a file (`/usr/include/stdio.h`) of standard routines and symbols that includes the definition of EOF.

If it is necessary to treat multiple files, you can use `cat` to collect the files for you:

```
tutorial% cat file1 file2 ... | ccstrip > output
```

and thus avoid learning how to access files from a program. By the way, the call to `exit` at the end is not necessary to make the program work properly, but it assures that any caller of the program will see a normal termination status (conventionally 0) from the program when it completes. Section 1.5.3 discusses returning status in more detail.

1.3. The Standard I/O Library

The 'Standard I/O Library' is a collection of routines intended to provide efficient and portable I/O services for most C programs. The standard I/O library is available on each system that supports C, so programs that confine their system interactions to its facilities can be transported from one system to another essentially without change.

This section discusses the basics of the standard I/O library. Section 1.7 — *The Standard I/O Library* — contains a more complete description of its capabilities and calling conventions.

Accessing Files

The above programs have all read the standard input and written the standard output, which we have assumed are magically predefined. The next step is to write a program that accesses a file that is *not* already connected to the program. One simple example is `wc`, which counts the lines, words and characters in a set of files. For instance, the command

```
tutorial% wc x.c y.c
```

displays the number of lines, words and characters in `x.c` and `y.c` and the totals.

The question is how to arrange for the named files to be read — that is, how to connect the filenames to the I/O statements which actually read the data.

The rules are simple — you have to *open* a file by the standard library function `fopen` before it can be read from or written to. `fopen` takes an external name (like `x.c` or `y.c`), does some housekeeping and negotiation with the operating system, and returns an internal name which must be used in subsequent reads or writes of the file.

This internal name is actually a pointer, called a *file pointer*, to a structure which contains information about the file, such as the location of a buffer, the current character position in the buffer, whether the file is being read or written, and the like. Users don't need to know the details, because part of the standard I/O definitions obtained by including `stdio.h` is a structure definition called `FILE`. The only declaration needed for a file pointer is exemplified by

```
FILE *fp, *fopen();
```

This says that `fp` is a pointer to a `FILE`, and `fopen` returns a pointer to a `FILE`. `FILE` is a type name, like `int`, not a structure tag.

The actual call to `fopen` in a program has the form:

```
fp = fopen(name, mode);
```

The first argument of `fopen` is the name of the file, as a character string. The second argument is the mode, also as a character string, which indicates how you intend to use the file. The allowable modes are read ("`r`"), write ("`w`"), or append ("`a`"). In addition, each *mode* may be followed by a `+` sign to open the file for reading and writing. "`r+`" positions the stream at the beginning of the file, "`w+`" creates or truncates the file, and "`a+`" positions the stream to the end of the file. Both reads and writes may be used on read/write streams, with the limitation that an `fseek`, `rewind`, or reading end-of-file must be used

between a read and a write or vice versa.

If a file that you open for writing or appending does not exist, it is created (if possible). Opening an existing file for writing discards the old contents. Trying to read a file that does not exist is an error, and there may be other causes of error as well (like trying to read a file when you don't have permission). If there is any error, `fopen` returns the null pointer value `NULL` — defined as zero in *stdio.h*.

The next thing needed is a way to read or write the file once it is open. There are several possibilities, of which `getc` and `putc` are the simplest. `getc` returns the next character from a file; it needs the file pointer to tell it what file. Thus

```
c = getc(fp)
```

places in `c` the next character from the file referred to by `fp`; it returns `EOF` when it reaches end of file. `putc` is the inverse of `getc`:

```
putc(c, fp)
```

puts the character `c` on the file `fp` and returns `c` as its value. `getc` and `putc` return `EOF` on error.

When a program is started, three streams are opened automatically, and file pointers are provided for them. These streams are the standard input, the standard output, and the standard error output; the corresponding file pointers are called `stdin`, `stdout`, and `stderr`. Normally these are all connected to the terminal, but may be redirected to files or pipes as described in Section 1.2. `stdin`, `stdout` and `stderr` are predefined in the I/O library as the standard input, output and error files; they may be used anywhere an object of type `FILE *` can be. They are constants, however, *not* variables, so don't try to assign to them.

With some of the preliminaries out of the way, we can now write `wc`. The basic design is one that has been found convenient for many programs: if there are command-line arguments, they are processed in order. If there are no arguments, the standard input is processed. This way the program can be used standalone or as part of a larger process.

```
#include <stdio.h>

main(argc, argv)    /* wc: count lines, words, chars */
int argc;
char *argv[ ];
{
    int c, i, inword;
    FILE *fp, *fopen();
    long linect, wordct, charct;
    long tlinect = 0, twordct = 0, tcharct = 0;

    i = 1;
    fp = stdin;
    do {
        if (argc > 1 && (fp=fopen(argv[i], "r")) == NULL) {
            fprintf(stderr, "wc: can't open %s\n", argv[i]);
            continue;
        }
        linect = wordct = charct = inword = 0;
        while ((c =getc(fp)) != EOF) {
            charct++;
            if (c == '\n')
                linect++;
            if (c == ' ' || c == '\t' || c == '\n')
                inword = 0;
            else if (inword == 0) {
                inword = 1;
                wordct++;
            }
        }
        printf("%7ld %7ld %7ld", linect, wordct, charct);
        printf(argc > 1 ? " %s\n" : "\n", argv[i]);
        fclose(fp);
        tlinect += linect;
        twordct += wordct;
        tcharct += charct;
    } while (++i < argc);
    if (argc > 2)
        printf("%7ld %7ld %7ld total\n", tlinect, twordct, tcharct);
    exit(0);
}
```

The function `fprintf` is identical to `printf`, save that the first argument is a file pointer that specifies the file to be written.

The function `fclose` is the inverse of `fopen`; it breaks the connection between the file pointer and the external name that was established by `fopen`, freeing the file pointer for another file. Since there is a limit on the number of files that a program may have open simultaneously, it's a good idea to free things when they are no longer needed. There is another reason to call `fclose` on an output file — it flushes the buffer in which `putc` is collecting output. `fclose` is called automatically for each open file when a program terminates normally.

Error Handling — Stderr and Exit

`stderr` is assigned to a program in the same way that `stdin` and `stdout` are. Output written on `stderr` appears on the user's terminal even if the standard output is redirected, unless the standard error is also redirected. `wc` writes its diagnostics on `stderr` instead of `stdout` so that if one of the files can't be accessed for some reason, the message finds its way to the user's terminal instead of disappearing down a pipeline or into an output file.

The argument of `exit` is made available to whatever process called the process that is exiting (see Section 1.5.3, so the success or failure of the program can be tested by another program that uses this one as a subprocess. By convention, a return value of 0 signals that all is well; nonzero values signal abnormal situations.

`exit` itself calls `fclose` for each open output file, to flush out any buffered output, then calls a routine named `_exit`. The function `_exit` terminates the program immediately without any buffer flushing; it may be called directly if desired.

Miscellaneous I/O Functions

The standard I/O library provides several other I/O functions besides those illustrated above.

Normally output with `putc`, and such is buffered — use `fflush(fp)` to force it out immediately.

`fscanf` is identical to `scanf`, except that its first argument is a file pointer (as with `fprintf`) that specifies the file from which the input comes; it returns EOF at end of file.

The functions `sscanf` and `sprintf` are identical to `fscanf` and `fprintf`, except that the first argument names a character string instead of a file pointer. The conversion is done from the string for `sscanf` and into it for `sprintf`, and no input or output is done.

`fgets(buf, size, fp)` copies the next line from `fp`, up to and including a newline, into `buf`; at most `size-1` characters are copied; it returns NULL at end of file. `fputs(buf, fp)` writes the string in `buf` onto file `fp`.

The function `ungetc(c, fp)` 'pushes back' the character `c` onto the input stream `fp`; a subsequent call to `getc`, `fscanf`, etc., will encounter `c`. Only one character of pushback per file is permitted.

1.4. Low-Level Input Output

This section describes the bottom level of I/O on the UNIX system. The lowest level of I/O in UNIX provides no buffering or any other services; it is in fact a direct entry into the operating system. You are entirely on your own, but on the other hand, you have the most control over what happens. And since the calls and usage are quite simple, this isn't as bad as it sounds.

File Descriptors

In the UNIX operating system, all input and output is done by reading or writing files, because all peripheral devices, even the user's terminal, are files in the file system. This means that a single, homogeneous interface handles all communication between a program and peripheral devices.

In the most general case, before reading or writing a file, it is necessary to inform the system of your intent to do so, a process called ‘opening’ the file. If you are going to write on a file, it may also be necessary to create it. The system checks your right to do so — does the file exist? Do you have permission to access it? — if all is well, returns a small positive integer called a *file descriptor*. Whenever I/O is to be done on the file, the file descriptor is used instead of the name to identify the file. This is roughly analogous to the use of `READ(5, . . .)` and `WRITE(6, . . .)` in FORTRAN. All information about an open file is maintained by the system; the user program refers to the file only by the file descriptor.

The file pointers discussed in Section 1.3 are similar in spirit to file descriptors, but file descriptors are more fundamental. A file pointer is a pointer to a structure that contains, among other things, the file descriptor for the file in question.

Since input and output involving the user’s terminal are so common, special arrangements exist to make this convenient. When the command interpreter (the ‘shell’) runs a program, it opens three files, with file descriptors 0, 1, and 2, called standard input, standard output, and standard error output. All of these are normally connected to the terminal, so if a program reads file descriptor 0 and writes file descriptors 1 and 2, it can do terminal I/O without opening the files.

If I/O is redirected to and from files with `<` and `>`, as in

```
tutorial% prog < infile > outfile
```

the shell changes the default assignments for file descriptors 0 and 1 from the terminal to the named files. Similar observations hold if the input or output is associated with a pipe. Normally file descriptor 2 remains attached to the terminal, so error messages can go there. In all cases, the file assignments are changed by the shell, not by the program. The program does not need to know where its input comes from nor where its output goes, so long as it uses file 0 for input and 1 and 2 for output.

read and write

All input and output is done by two functions called `read` and `write`. For both, the first argument is a file descriptor. The second argument is a buffer in your program where the data is to come from or go to. The third argument is the number of bytes to be transferred. The calls are

```
n_read = read(fd, buf, n);
```

```
n_written = write(fd, buf, n);
```

Each call returns a byte count which is the number of bytes actually transferred. On reading, the number of bytes returned may be less than the number asked for, because fewer than `n` bytes remained to be read. When the file is a terminal, `read` normally reads only up to the next newline, which is generally less than what was requested. A return value of zero bytes implies end of file, and `-1` indicates an error of some sort. For writing, the returned value is the number of bytes actually written; it is generally an error if this isn’t equal to the number supposed to be written.

The number of bytes to be read or written is quite arbitrary. The two most common values are 1, which means one character at a time ('unbuffered'), and 1024, corresponding to a physical blocksize on many peripheral devices. This latter size will be most efficient, but even character-at-a-time I/O is not inordinately expensive.

Putting these facts together, we can write a simple program to copy its input to its output. This program will copy anything to anything, since the input and output can be redirected to any file or device.

```
#define BUFSIZE 1024

main() /* copy input to output */
{
    char    buf[BUFSIZE];
    int n;

    while ((n = read(0, buf, BUFSIZE)) > 0)
        write(1, buf, n);
    exit(0);
}
```

If the file size is not a multiple of BUFSIZE, some read will return a smaller number of bytes, and the next call to read after that will return zero.

It is instructive to see how read and write can be used to construct higher-level routines like getchar, putchar, etc. For example, here is a version of getchar which does unbuffered input.

```
#define CMASK    0377    /* for making char's > 0 */

getchar() /* unbuffered single character input */
{
    char c;

    return((read(0, &c, 1) > 0) ? c & CMASK : EOF);
}
```

c *must* be declared char, because read accepts a character pointer. The character being returned must be masked with 0377 to ensure that it is positive; otherwise sign extension may make it negative. The constant 0377 is appropriate for the Sun but not necessarily for other machines.

The second version of getchar does input in big chunks, and hands out the characters one at a time:

```
#define CMASK 0377 /* for making char's > 0 */
#define BUFSIZE 1024

getchar() /* buffered version */
{
    static char buf[BUFSIZE];
    static char *bufp = buf;
    static int n = 0;

    if (n == 0) { /* buffer is empty */
        n = read(0, buf, BUFSIZE);
        bufp = buf;
    }
    return((--n >= 0) ? *bufp++ & CMASK : EOF);
}
```

Open, Creat, Close, Unlink

Other than the default standard input, output and error files, you must explicitly open files in order to read or write them. There are two system entry points for this, `open` and `creat`.

`open` is rather like the `fopen` discussed in the previous section, except that instead of returning a file pointer, it returns a file descriptor, which is just an `int`.

```
int fd;

fd = open(name, rmode);
```

As with `fopen`, the `name` argument is a character string corresponding to the external file name. The access mode argument is different, however: `rmode` is 0 for read, 1 for write, and 2 for read and write access. `open` returns `-1` if any error occurs; otherwise it returns a valid file descriptor.

It is an error to try to open a file that does not exist. The entry point `creat` is provided to create new files, or to rewrite old ones.

```
fd = creat(name, pmode);
```

returns a file descriptor if it could create the file called `name`, and `-1` if not. If the file already exists, `creat` will truncate it to zero length; it is not an error to create a file that already exists.

If the file is brand new, `creat` creates it with the *protection mode* specified by the `pmode` argument. In the UNIX file system, there are nine bits of protection information associated with a file, controlling read, write and execute permission for the owner of the file, for the owner's group, and for all others. Thus a three-digit octal number is most convenient for specifying the permissions. For example, `0755` specifies read, write and execute permission for the owner, and read and execute permission for the group and everyone else.

To illustrate, here is a simplified version of the UNIX utility `cp`, a program which copies one file to another. The main simplification is that our version copies only one file, and does not permit the second argument to be a directory:


```

#define NULL 0
#define BUFSIZE 1024
#define PMODE 0644 /* RW for owner, R for group, others */

main(argc, argv)    /* cp: copy f1 to f2 */
int argc;
char *argv[ ];
{
    int f1, f2, n;
    char    buf[BUFSIZE];

    if (argc != 3)
        error("Usage: cp from to", NULL);
    if ((f1 = open(argv[1], 0)) == -1)
        error("cp: can't open %s", argv[1]);
    if ((f2 = creat(argv[2], PMODE)) == -1)
        error("cp: can't create %s", argv[2]);

    while ((n = read(f1, buf, BUFSIZE)) > 0)
        if (write(f2, buf, n) != n)
            error("cp: write error", NULL);
    exit(0);
}

error(s1, s2)    /* print error message and die */
char *s1, *s2;
{
    printf(s1, s2);
    printf("\n");
    exit(1);
}

```

As we said earlier, there is a limit (typically 20-32) on the number of files which a program may have open simultaneously. Accordingly, any program which intends to process many files must be prepared to reuse file descriptors. The routine `close` breaks the connection between a file descriptor and an open file, and frees the file descriptor for use with some other file. Termination of a program via `exit` or return from the main program closes all open files.

The function `unlink(filename)` removes the file `filename` from the file system.

Random Access — Seek and Lseek

File I/O is normally sequential: each `read` or `write` takes place at a position in the file right after the previous one. When necessary, however, a file can be read or written in any arbitrary order. The system call `lseek` provides a way to move around in a file without actually reading or writing:

```
lseek(fd, offset, origin);
```

forces the current position in the file whose descriptor is `fd` to move to position `offset`, which is taken relative to the location specified by `origin`. Subsequent reading or writing will begin at that position. `offset` is a `long`; `fd` and `origin` are `int`'s. `origin` can be 0, 1, or 2 to specify that `offset` is to be

measured from the beginning, from the current position, or from the end of the file, respectively. For example, to append to a file, seek to the end before writing:

```
lseek(fd, 0L, 2);
```

To get back to the beginning ('rewind'),

```
lseek(fd, 0L, 0);
```

Notice the 0L argument; it could also be written as (long) 0.

With `lseek`, it is possible to treat files more or less like large arrays, at the price of slower access. For example, the following simple function reads any number of bytes from any arbitrary place in a file.

```
get(fd, pos, buf, n) /* read n bytes from position pos */
int fd, n;
long pos;
char *buf;
{
    lseek(fd, pos, 0); /* get to pos */
    return(read(fd, buf, n));
}
```

Error Processing

The routines discussed in this section, and in fact all the routines which are direct entries into the system can incur errors. Usually they indicate an error by returning a value of `-1`. Sometimes it is nice to know what sort of error occurred; for this purpose all these routines, when appropriate, leave an error number in the external variable `errno`. The meanings of the various error numbers are listed in *intro(2)* in the Sun *UNIX Interface Reference Manual* so your program can, for example, determine if an attempt to open a file failed because it did not exist or because the user lacked permission to read it. Perhaps more commonly, you may want to display the reason for failure. The routine `perror` displays a message associated with the value of `errno`; more generally, `sys_errno` is an array of character strings which can be indexed by `errno` and displayed by your program.

1.5. Processes

It is often easier to use a program written by someone else than to invent one's own. This section describes how to execute a program from within another.

The 'System' Function

The easiest way to execute a program from another is to use the standard library routine `system`. `system` takes one argument, a command string exactly as typed at the terminal (except for the newline at the end) and executes it. For instance, to timestamp the output of a program,

```
main( ) {
    system("date"); /* rest of processing */
}
```

If the command string has to be built from pieces, the in-memory formatting capabilities of `sprintf` may be useful.

Low-Level Process Creation — Exec1 and Execv

Remember that `getc` and `putc` normally buffer their input; terminal I/O will not be properly synchronized unless this buffering is defeated. For output, use `fflush`; for input, see `setbuf` in section 1.7.

If you're not using the standard library, or if you need finer control over what happens, you will have to construct calls to other programs using the more primitive routines that the standard library's `system` routine is based on¹.

The most basic operation is to execute another program *without returning*, by using the routine `exec1`. To display the date as the last action of a running program, use

```
exec1("/bin/date", "date", NULL);
```

The first argument to `exec1` is the *filename* of the command; you have to know where it is found in the file system. The second argument is conventionally the program name (that is, the last component of the file name), but this is seldom used except as a placeholder. If the command takes arguments, they are strung out after this; the end of the list is marked by a `NULL` argument.

The `exec1` call overlays the existing program with the new one, runs that, then exits. There is *no* return to the original program.

More realistically, a program might fall into two or more phases that communicate only through temporary files. Here it is natural to start the second pass simply by an `exec1` call from the first.

The one exception to the rule that the original program never gets control back occurs when there is an error, for example if the file can't be found or is not executable. If you don't know where `date` is located, you might try

```
exec1("/bin/date", "date", NULL);
exec1("/usr/bin/date", "date", NULL);
fprintf(stderr, "Someone stole 'date'\n");
```

A variant of `exec1` called `execv` is useful when you don't know in advance how many arguments there are going to be. The call is

```
execv(filename, argp);
```

where `argp` is an array of pointers to the arguments; the last pointer in the array must be `NULL` so `execv` can tell where the list ends. As with `exec1`, `filename` is the file in which the program is found, and `argp[0]` is the name of the program. (This arrangement is identical to the `argv` array for program arguments.)

Neither of these routines provides the niceties of normal command execution. There is no automatic search of multiple directories — you have to know precisely where the command is located. Nor do you get the expansion of metacharacters like `<`, `>`, `*`, `?`, and `[]` in the argument list. If you want these, use `exec1` to invoke the shell `sh`, which then does all the work. Construct a string

¹ `system` uses `/bin/sh` (the Bourne Shell) to execute the command string, so syntax specific to the C-Shell will not work.

commandline that contains the complete command as it would have been typed at the terminal, then say

```
execl("/bin/sh", "sh", "-c", commandline, NULL);
```

The shell is assumed to be at a fixed place, `/bin/sh`. Its argument `-c` says to treat the next argument as a whole command line, so it does just what you want. The only problem is in constructing the right information in `commandline`.

Control of Processes — Fork and Wait

So far what we've talked about isn't really all that useful by itself. Now we will show how to regain control after running a program with `execl` or `execv`. Since these routines simply overlay the new program on the old one, to save the old one requires that it first be split into two copies; one of these can be overlaid, while the other waits for the new, overlaying program to finish. The splitting is done by a routine called `fork`:

```
proc_id = fork( );
```

splits the program into two copies, both of which continue to run. The only difference between the two is the value of `proc_id`, the 'process id.' In one of these processes (the 'child'), `proc_id` is zero. In the other (the 'parent'), `proc_id` is nonzero; it is the process number of the child. Thus the basic way to call, and return from, another program is

```
if (fork() == 0)
    execl("/bin/sh", "sh", "-c", cmd, NULL); /* in child */
```

And in fact, except for handling errors, this is sufficient. The `fork` makes two copies of the program. In the child, the value returned by `fork` is zero, so it calls `execl` which does the command and then dies. In the parent, `fork` returns nonzero so it skips the `execl`. If there is any error, `fork` returns `-1`.

More often, the parent wants to wait for the child to terminate before continuing itself. This can be done with the function `wait`:

```
int status;

if (fork( ) == 0)
    execl(...);
wait(&status);
```

This still doesn't handle any abnormal conditions, such as a failure of the `execl` or `fork`, or the possibility that there might be more than one child running simultaneously. The `wait` returns the process id of the terminated child, if you want to check it against the value returned by `fork`. Finally, this fragment doesn't deal with any funny behavior on the part of the child (which is reported in `status`). Still, these three lines are the heart of the standard library's `system` routine, which we'll show in a moment.

The `status` returned by `wait` encodes in its low-order eight bits the system's idea of the child's termination status; it is 0 for normal termination and nonzero to indicate various kinds of problems. The next higher eight bits are taken from the argument of the call to `exit` which caused a normal termination of the child process. It is good coding practice for all programs to return meaningful status.

When a program is called by the shell, the three file descriptors 0, 1, and 2 are set up to point at the right files (see Section 1.4.1), and all other possible file descriptors are available for use. When this program calls another one, correct etiquette suggests making sure the same conditions hold. Neither `fork` nor the `exec` calls affects open files in any way. If the parent is buffering output that must come out before output from the child, the parent must flush its buffers before the `exec1`. Conversely, if a caller buffers an input stream, the called program will lose any information that has been read by the caller.

Pipes

A *pipe* is an I/O channel intended for use between two cooperating processes: one process writes into the pipe, while the other process reads from the pipe. The system looks after buffering the data and synchronizing the two processes. Most pipes are created by the shell, as in

```
tutorial% ls | pr
```

which connects the standard output of `ls` to the standard input of `pr`. Sometimes, however, it is most convenient for a process to set up its own plumbing; in this section, we illustrate how the pipe connection is established and used.

The system call `pipe` creates a pipe. Since a pipe is used for both reading and writing, two file descriptors are returned; the actual usage is like this:

```
int    fd[2];

stat = pipe(fd);
if (stat == -1)
    /* there was an error ... */
```

`fd` is an array of two file descriptors, where `fd[0]` is the read side of the pipe and `fd[1]` is for writing. These may be used in `read`, `write` and `close` calls just like any other file descriptors.

If a process reads a pipe which is empty, it waits until data arrives; if a process writes into a pipe which is too full, it waits until the pipe empties somewhat. If the write side of the pipe is closed, a subsequent `read` will encounter end of file.

To illustrate the use of pipes in a realistic setting, let us write a function called `popen(cmd, mode)`, which creates a process `cmd` (just as `system` does), and returns a file descriptor that will either read or write that process, according to `mode`. That is, the call

```
fout = popen("pr", WRITE);
```

creates a process that executes the `pr` command; subsequent `write` calls using the file descriptor `fout` will send their data to that process through the pipe.

`popen` first creates the pipe with a `pipe` system call; it then `fork`'s to create two copies of itself. The child decides whether it is supposed to read or write, closes the other side of the pipe, then calls the shell (via `exec1`) to run the desired process. The parent likewise closes the end of the pipe it does not use. These closes are necessary to make end-of-file tests work properly. For example, if a child that intends to read fails to close the write end of the pipe, it will never

see the end of the pipe file, just because there is one writer potentially active.

```
#include <stdio.h>

#define READ    0
#define WRITE   1
#define tst(a, b)  (mode == READ ? (b) : (a))
static int popen_pid;

popen(cmd, mode)
char *cmd;
int mode;
{
    int p[2];

    if (pipe(p) < 0)
        return(NULL);
    if ((popen_pid = fork( )) == 0) {
        close(tst(p[WRITE], p[READ]));
        close(tst(0, 1));
        dup(tst(p[READ], p[WRITE]));
        close(tst(p[READ], p[WRITE]));
        execl("/bin/sh", "sh", "-c", cmd, 0);
        _exit(1); /* disaster has occurred if we get here */
    }
    if (popen_pid == -1)
        return(NULL);
    close(tst(p[READ], p[WRITE]));
    return(tst(p[WRITE], p[READ]));
}
```

The sequence of `close`'s in the child is a bit tricky. Suppose that the task is to create a child process that will read data from the parent. Then the first `close` closes the write side of the pipe, leaving the read side open. The lines

```
close(tst(0, 1));
dup(tst(p[READ], p[WRITE]));
```

are the conventional way to associate the pipe descriptor with the standard input of the child. The `close` closes file descriptor 0, that is, the standard input. `dup` is a system call that returns a duplicate of an already open file descriptor. File descriptors are assigned in increasing order and the first available one is returned, so the effect of the `dup` is to copy the file descriptor for the pipe (read side) to file descriptor 0; thus the read side of the pipe becomes the standard input². Finally, the old read side of the pipe is closed.

A similar sequence of operations takes place when the child process is supposed to write to the parent instead of reading. You may find it a useful exercise to step through that case.

² Yes, this is a bit tricky, but it's a standard idiom.

The job is not quite done, for we still need a function `pclose` to close the pipe created by `popen`. The main reason for using a separate function rather than `close` is that it is desirable to wait for the termination of the child process. First, the return value from `pclose` indicates whether the process succeeded. Equally important when a process creates several children is that only a bounded number of unwaited-for children can exist, even if some of them have terminated; performing the `wait` lays the child to rest. Thus:

```
#include <signal.h>

pclose(fd) /* close pipe fd */
int fd;
{
    register r, (*hstat)(), (*istat)(), (*qstat)();
    int status;
    extern int popen_pid;

    close(fd);
    istat = signal(SIGINT, SIG_IGN);
    qstat = signal(SIGQUIT, SIG_IGN);
    hstat = signal(SIGHUP, SIG_IGN);
    while ((r = wait(&status)) != popen_pid && r != -1);
    if (r == -1)
        status = -1;
    signal(SIGINT, istat);
    signal(SIGQUIT, qstat);
    signal(SIGHUP, hstat);
    return(status);
}
```

The calls to `signal` make sure that no interrupts, etc. interfere with the waiting process; this is the topic of the next section.

The routine as written has the limitation that only one pipe may be open at once, because of the single shared variable `popen_pid`; it really should be an array indexed by file descriptor. A `popen` function, with slightly different arguments and return value is available as part of the standard I/O library discussed below. As currently written, it shares the same limitation.

1.6. Signals — Interrupts and All That

This section is concerned with how to deal gracefully with signals from the outside world (like interrupts), and with program faults. Since there's nothing very useful that can be done from within C about program faults, which arise mainly from illegal memory references or from execution of peculiar instructions, we'll discuss only the outside world signals: *interrupt* and *quit*, which are generated from the keyboard³, *hangup*, caused by hanging up the phone on dialup lines, and *terminate*, generated by the *kill* command. When one of these events occurs, the signal is sent to *all* processes which were started from the corresponding terminal — the signal terminates the process unless other arrangements have been

³ The current binding of characters and signals can be discovered by the `stty all` command. On Sun systems, typing control-C usually generates the *kill* signal and control-\ generates the *quit* signal.

made. In the *quit* case, a core image file is written for debugging purposes.

`signal` is the routine which alters the default action. `signal` has two arguments: the first specifies the signal to be processed, and the second argument specifies what to do with that signal. The first argument is just a numeric code, but the second is either a function, or a somewhat strange code that requests that the signal either be ignored or that it be given the default action. The include file *signal.h* gives names for the various arguments, and should always be included when signals are used. Thus

```
#include <signal.h>
...
signal(SIGINT, SIG_IGN);
```

means that interrupts are ignored, while

```
signal(SIGINT, SIG_DFL);
```

restores the default action of process termination. In all cases, `signal` returns the previous value of the signal. The second argument to `signal` may instead be the name of a function (which has to be declared explicitly if the compiler hasn't seen it already). In this case, the named routine will be called when the signal occurs. Most commonly this facility is used so that the program can clean up unfinished business before terminating, for example to delete a temporary file:

```
#include <signal.h>

main( )
{
    int onintr( );

    if (signal(SIGINT, SIG_IGN) != SIG_IGN)
        signal(SIGINT, onintr);

    /* Process ... */

    exit(0);
}

onintr( )
{
    unlink(tempfile);
    exit(1);
}
```

Why the test and the double call to `signal`? Recall that signals like `interrupt` are sent to *all* processes started from a particular terminal. Accordingly, when a program is to be run non-interactively (started by `&`), the shell turns off interrupts for it so it won't be stopped by interrupts intended for foreground processes. If this program began by announcing that all interrupts were to be sent to the `onintr` routine regardless, that would undo the shell's effort to protect it when run in the background.

The solution, shown above, is to test the state of interrupt handling, and to continue to ignore interrupts if they are already being ignored. The code as written

depends on the fact that `signal` returns the previous state of a particular signal. If signals were already being ignored, the process should continue to ignore them; otherwise, they should be caught.

A more sophisticated program may wish to intercept an interrupt and interpret it as a request to stop what it is doing and return to its own command processing loop. Think of a text editor: interrupting a long display should not terminate the edit session and lose the work already done. The outline of the code for this case is probably best written like this:

```
#include <signal.h>
#include <setjmp.h>
jmp_buf sjbuf;

main( )
{
    int (*istat)( ), onintr( );

    istat = signal(SIGINT, SIG_IGN); /* original status */
    setjmp(sjbuf); /* save current stack position */
    if (istat != SIG_IGN)
        signal(SIGINT, onintr);

    /* main processing loop */
}

onintr( )
{
    printf("\nInterrupt\n");
    longjmp(sjbuf); /* return to saved state */
}
```

The include file *setjmp.h* declares the type `jmp_buf` — an object in which the state can be saved. `sjbuf` is such an object. The `setjmp` routine then saves the state of things. When an interrupt occurs the `onintr` routine is called, which can display a message, set flags, or whatever. `longjmp` takes as argument an object set by `setjmp`, and restores control to the location following the call to `setjmp`, so control (and the stack level) will pop back to the place in the main routine where the signal is set up and the main loop entered. Notice, by the way, that the signal gets set again after an interrupt occurs. This is necessary; most signals are automatically reset to their default action when they occur.

Some programs that want to detect signals simply can't be stopped at an arbitrary point, for example in the middle of updating a linked list. If the routine called when a signal occurs sets a flag and then returns instead of calling `exit` or `longjmp`, execution continues at the exact point it was interrupted. The interrupt flag can then be tested later.

There is one difficulty associated with this approach. Suppose the program is reading the terminal when the interrupt is sent. The specified routine is duly called; it sets its flag and returns. If it were really true, as we said above, that 'execution resumes at the exact point it was interrupted,' the program would continue reading the terminal until the user typed another line. This behavior might

well be confusing, since the user might not know that the program is reading; he presumably would prefer to have the signal take effect instantly. The method chosen to resolve this difficulty is to terminate the terminal read when execution resumes after the signal, returning an error code which indicates what happened.

Thus programs which catch and resume execution after signals should be prepared for 'errors' which are caused by interrupted system calls. The ones to watch out for are reads from a terminal, `wait`, and `pause`. A program whose `onintr` routine just sets `intflag`, resets the interrupt signal, and returns, should usually include code like the following when it reads the standard input:

```
if (getchar() == EOF)
    if (intflag)
        /* EOF caused by interrupt */
    else
        /* true end-of-file */
```

A final subtlety to keep in mind becomes important when catching signals is combined with executing other programs. Suppose a program catches interrupts, and also includes a method (like '!' in the editor) whereby other programs can be executed. Then the code should look something like this:

```
if (fork( ) == 0)
    execl(...);
signal(SIGINT, SIG_IGN); /* ignore interrupts */
wait(&status);          /* until the child is done */
signal(SIGINT, onintr); /* restore interrupts */
```

Why is this? Again, it's not obvious, but not really difficult. Suppose the program you call catches its own interrupts. If you interrupt the subprogram, it will get the signal and return to its main loop, and probably read your terminal. But the calling program will also pop out of its wait for the subprogram and read your terminal. Having two processes reading your terminal is very unfortunate, since the system figuratively flips a coin to decide who should get each line of input. A simple way out is to have the parent program ignore interrupts until the child is done. This reasoning is reflected in the standard I/O library function `system`:

```

#include <signal.h>

system(s) /* run command string s */
char *s;
{
    int status, pid, w;
    register int (*istat)(), (*qstat)();

    if ((pid = fork()) == 0) {
        execl("/bin/sh", "sh", "-c", s, 0);
        _exit(127);
    }
    istat = signal(SIGINT, SIG_IGN);
    qstat = signal(SIGQUIT, SIG_IGN);
    while ((w = wait(&status)) != pid && w != -1)
        ;
    if (w == -1)
        status = -1;
    signal(SIGINT, istat);
    signal(SIGQUIT, qstat);
    return(status);
}

```

As an aside on declarations, the function `signal` obviously has a rather strange second argument. It is in fact a pointer to a function delivering an integer, and this is also the type of the signal routine itself. The two values `SIG_IGN` and `SIG_DFL` have the right type, but are chosen so they coincide with no possible actual functions. For the enthusiast, here is how they are defined for the Sun system — the definitions should be sufficiently ugly and nonportable to encourage use of the include file.

```

#define SIG_DFL      (int (*)())0
#define SIG_IGN     (int (*)())1

```

1.7. The Standard I/O Library

The standard I/O library was designed with the following goals in mind:

1. It must be as efficient as possible, both in time and in space, so that there will be no hesitation in using it, no matter how critical the application.
2. It must be simple to use, and also free of the magic numbers and mysterious calls whose use mars the understandability and portability of many programs using older packages.
3. The interface provided should be applicable on all machines, whether or not the programs which implement it are directly portable to other systems, or to machines non-Sun running a version of UNIX.

General Usage

Each program using the library must have the line

```
#include <stdio.h>
```

which defines certain macros and variables. The routines are in the normal C library, so no special library argument is needed for loading. All names in the

include file intended only for internal use begin with an underscore `_` to reduce the possibility of collision with a user name. The names intended to be visible outside the package are

`stdin` the name of the standard input stream
`stdout` the name of the standard output stream
`stderr` the name of the standard error stream
`EOF` is actually `-1`, and is the value returned by the read routines on end-of-file or error
`NULL` is a notation for the null pointer, returned by pointer-valued functions to indicate an error
`FILE` expands to `struct _iob` and is a useful shorthand when declaring pointers to streams
`BUFSIZ` is a number (viz. 1024) of the size suitable for an I/O buffer supplied by the user. See `setbuf`, below
`getc`, `getchar`, `putc`, `putchar`, `feof`, `ferror`, `fileno` are defined as macros. Their actions are described below; they are mentioned here to point out that it is not possible to redeclare them and that they are not actually functions; thus, for example, they may not have breakpoints set on them.

The routines in this package offer the convenience of automatic buffer allocation and output flushing where appropriate. The names `stdin`, `stdout`, and `stderr` are constants and may not be assigned to.

Standard I/O Library Calls

```
FILE *fopen(filename, type)
char *filename;
char *type;
```

opens the file and, if needed, allocates a buffer for it. `filename` is a character string specifying the name. `type` is a character string (not a single character). It may be `"r"`, `"w"`, or `"a"` to indicate intent to read, write, or append. In addition, each *mode* may be followed by a `+` sign to open the file for reading and writing. `"r+"` positions the stream at the beginning of the file, `"w+"` creates or truncates the file, and `"a+"` positions the stream to the end of the file. Both reads and writes may be used on read/write streams, with the limitation that an `fseek`, `rewind`, or reading end-of-file must be used between a read and a write or vice versa. The value returned is a file pointer. If it is `NULL` the attempt to open failed.

`freopen`

```
FILE *freopen(filename, type, ioptr)
char *filename;
char *type;
FILE *ioptr;
```

The stream named by `ioptr` is closed, if necessary, and then reopened as if by `fopen`. If the attempt to open fails, `NULL` is returned, otherwise `ioptr` is returned, which now refers to the new file. Often the reopened stream is `stdin`

or `stdout`. The filename and type parameters are as for `fopen`.

```
getc          int getc(ioptr)
              FILE *ioptr;
```

returns the next character from the stream named by `ioptr`, which is a pointer to a file such as returned by `fopen`, or the name `stdin`. The integer EOF is returned on end-of-file or when an error occurs. The null character `\0` is a legal character.

```
fgetc        int fgetc(ioptr)
              FILE *ioptr;
```

acts like `getc` but is a genuine function, not a macro, so it can be pointed to, passed as an argument, etc.

```
putc         int putc(c, ioptr)
              int c;
              FILE *ioptr;
```

`putc` writes the character `c` on the output stream named by `ioptr`, which is a value returned from `fopen` or perhaps `stdout` or `stderr`. The character is returned as value, and EOF is returned on error.

```
fputc        int fputc(c, ioptr)
              int c;
              FILE *ioptr;
```

acts like `putc` but is a genuine function, not a macro.

```
fclose       int fclose(ioptr)
              FILE *ioptr;
```

The file corresponding to `ioptr` is closed after any buffers are emptied. A buffer allocated by the I/O system is freed. `fclose` is automatic on normal termination of the program.

```
fflush       int fflush(ioptr)
              FILE *ioptr;
```

Any buffered information on the (output) stream named by `ioptr` is written out. Output files are normally buffered if they are not directed to the terminal.

```
exit         (void) exit(errcode);
              int errcode;
```

terminates the process and returns its argument as status to the parent. This is a special version of the routine which calls `fflush` for each output file. To terminate without flushing, use `_exit`.

```
feof         int feof(ioptr)
              FILE *ioptr;
```

returns nonzero when end-of-file has occurred on the specified input stream.

- ferror** `int ferror(ioptr)`
 `FILE *ioptr;`
- returns nonzero when an error has occurred while reading or writing the named stream. The error indication lasts until the file has been closed.
- getchar** `int getchar();`
- is identical to `getc(stdin)`.
- putchar** `int putchar(c);`
- is identical to `putc(c, stdout)`.
- fgets** `char *fgets(s, n, ioptr)`
 `char *s;`
 `int n;`
 `FILE *ioptr;`
- reads to `n-1` characters, or up to a newline character, whichever comes first, from the stream `ioptr` into the string pointed to by the character pointer `s`. A null character is placed after the last character read in the strings `s`. `fgets` returns the first argument, or `NULL` if error or end-of-file occurred.
- puts** `int puts(s)`
 `char *s;`
- `puts` copies the null-terminated strings specified by `s` onto the standard output stream and appends a newline character.
- fputs** `int fputs(s, ioptr)`
 `char *s;`
 `FILE *ioptr;`
- writes the null-terminated string (character array) `s` on the stream `ioptr`. No newline is appended. The last character transmitted is returned as value, or `EOF` is returned on error.
- ungetc** `int ungetc(c, ioptr)`
 `int c;`
 `FILE *ioptr;`
- The argument character `c` is pushed back on the input stream named by `ioptr`. Only one character may be pushed back.
- printf**

```

int printf(format, a1, ...)
    char *format;

int fprintf(ioptr, format, a1, ...)
    FILE *ioptr;
    char *format;

int sprintf(s, format, a1, ...)
    char *s;
    char *format;

```

`printf` writes on the standard output. `fprintf` writes on the output stream named by `ioptr`. `sprintf` puts characters in the character array (string) named by `s`. The specifications are as described in `printf(3)` in the Sun *UNIX Interface Reference Manual*.

`printf` and `fprintf` return the number of characters actually transmitted, or return EOF if any error condition exists on the output file. `sprintf` returns a pointer to the buffer where the formatted string is placed.

scanf

```

int scanf(format, a1, ...)
    char *format;

int fscanf(ioptr, format, a1, ...)
    FILE *ioptr;
    char *format;

int sscanf(s, format, a1, ...)
    char *s;
    char *format;

```

`scanf` reads from the standard input. `fscanf` reads from the named input stream. `sscanf` reads from the character string supplied as `s`. `scanf` reads characters, interprets them according to the `format`, and stores the results in its arguments. Each routine expects as arguments a control string `format`, and a set of arguments, *each of which must be a pointer*, indicating where the converted input should be stored.

`scanf` returns as its value the number of successfully matched and assigned input items. This can be used to decide how many input items were found. On end of file, EOF is returned; note that this is different from 0, which means that the next input character does not match what was called for in the control string.

fread

```

int fread(ptr, sizeof(*ptr), nitems, ioptr)
    unsigned nitems;
    FILE *ioptr;

```

reads `nitems` of data of the type of `*ptr` from file `ioptr` into the memory area starting at `ptr`. No advance notification that binary I/O is being done is required. `fread` returns the number of items actually read from the specified stream.

`fwrite` `int fwrite(ptr, sizeof(*ptr), nitems, ioptr)`
 `unsigned nitems;`
 `FILE *ioptr;`

Like `fread`, but in the other direction. `fwrite` returns the number of items actually transmitted to the specified stream. This may possibly be less than the number of items requested if an error occurs while the transfer is in process.

`rewind` `(void) rewind(ioptr)`
 `FILE *ioptr;`

rewinds the stream named by `ioptr`. It is not very useful except on input, since a rewound output file is still open only for output.

`system` `int system(string)`
 `char *string;`

The `string` is executed by the shell as if typed at the terminal. The return value is the exit code of the invoked shell, which is usually the exit code of the last command executed by it.

`getw` `int getw(ioptr)`
 `FILE *ioptr;`

returns the next word from the input stream named by `ioptr`. EOF is returned on end-of-file or error, but since this a perfectly good integer, `feof` and `ferror` should be used. A 'word' is 32 bits on the Sun Workstation.

`putw` `int putw(w, ioptr)`
 `FILE *ioptr;`

writes the integer `w` on the named output stream. `putw` returns the current error status of the specified stream, as if an `ferror` call had been made.

`setbuf` `(void) setbuf(ioptr, buf)`
 `FILE *ioptr; char *buf;`

`setbuf` may be used after a stream has been opened but before I/O has started. If `buf` is NULL, the stream is unbuffered. Otherwise the buffer supplied is used. It must be a character array of sufficient size:

```
char    buf[BUFSIZ];
```

`setbuffer` `(void) setbuffer(ioptr, buf, size)`
 `FILE *ioptr;`
 `char *buf;`
 `int size;`

`setbuffer` is like `setbuf` (described above), but can be used when a specified, nonstandard buffer size should be used.

`fileno` `int fileno(ioptr)`
 `FILE *ioptr;`

returns the integer file descriptor associated with the file.

```
fseek      int fseek(ioptr, offset, ptrname)
           FILE *ioptr;
           long offset;
           int ptrname;
```

The location of the next byte in the stream named by `ioptr` is adjusted. `offset` is a long integer. If `ptrname` is 0, the offset is measured from the beginning of the file; if `ptrname` is 1, the offset is measured from the current read or write pointer; if `ptrname` is 2, the offset is measured from the end of the file. The routine accounts properly for any buffering. When this routine is used on non UNIX systems, the offset must be a value returned from `ftell` and the `ptrname` must be 0.

```
ftell      long ftell(ioptr)
           FILE *ioptr;
```

The byte offset, measured from the beginning of the file, associated with the named stream is returned. Any buffering is properly accounted for. On non UNIX systems the value of this call is useful only for handing to `fseek`, so as to position the file to the same place it was when `ftell` was called.

```
getpw      int getpw(uid, buf)
           int uid;
           char *buf;
```

The password file is searched for the given integer user ID. If an appropriate line is found, it is copied into the character array `buf`, and 0 is returned. If no line is found corresponding to the user ID then 1 is returned.

```
malloc      char *malloc(num)
           int num;
```

allocates `num` bytes. The pointer returned is aligned so as to be usable for any purpose. `NULL` is returned if no space is available.

```
free        int free(ptr)
           char *ptr;
```

`free` frees up memory previously allocated by `malloc`. `free` returns a 0 if any errors were detected (such as `ptr` being misaligned), and returns 1 otherwise. Disorder can be expected if the pointer was not obtained from `malloc`.

```
calloc      char *calloc(num, size);
           unsigned num;
           unsigned size;
```

allocates space for `num` items, each of size `size`. The space is guaranteed to be set to 0 and the pointer is aligned so as to be usable for any purpose. `NULL` is returned if no space is available.

```
cfree          (void) cfree(ptr, num, size)
               char *ptr;
               unsigned num;
               unsigned size;
```

Space is returned to the pool used by `calloc`. Disorder can be expected if the pointer was not obtained from `calloc`.

The following are macros whose definitions may be obtained by including `<ctype.h>`.

Character Type Checking

`isalpha(c)` returns nonzero if *c* is alphabetic.

`isupper(c)` returns nonzero if *c* is upper-case alphabetic.

`islower(c)` returns nonzero if *c* is lower-case alphabetic.

`isdigit(c)` returns nonzero if *c* is a digit.

`isxdigit(c)` returns nonzero if *c* is a hexadecimal digit — that is, one of '0' through '9', 'a' through 'f', or 'A' through 'F'.

`isspace(c)` returns nonzero if *c* is a spacing character: tab, newline, carriage return, vertical tab, form feed, space.

`ispunct(c)` returns nonzero if *c* is any punctuation character, that is, not a space, letter, digit or control character.

`isalnum(c)` returns nonzero if *c* is a letter or a digit.

`isprint(c)` returns nonzero if *c* is printable — a letter, digit, space, or punctuation character.

`iscntrl(c)` returns nonzero if *c* is a control character.

`isascii(c)` returns nonzero if *c* is an ASCII character, that is, less than octal 0200.

`isgraph(c)` returns nonzero if *c* is a printing character — like `isprint(c)` but doesn't include the space character.

Character Type Conversion

`toupper(c)` returns the upper-case character corresponding to the lower-case letter *c*.

`tolower(c)` returns the lower-case character corresponding to the upper-case letter *c*.

Tools for the C Programming Language

Tools for the C Programming Language	33
2.1. <code>ctags</code> — Build Index File for C Functions	33
2.2. <code>Lint</code> — A C Program Checker	35
Using <code>Lint</code>	35
A Word About Philosophy	35
Unused Variables and Functions	36
Set/Used Information	36
Flow of Control	37
Function Values	37
Type Checking	38
Type Casts	39
Nonportable Character Use	39
Assignments of longs to ints	39
Strange Constructions	40
Ancient History	40
Pointer Alignment	41
Multiple Uses and Side Effects	41
Implementation	42
Portability	42
Shutting <code>Lint</code> Up	44
Library Declaration Files	45
Bugs, etc.	45
Current <code>Lint</code> Options	46

Tools for the C Programming Language

Utilities described in this chapter cover facilities for the C programming language.

- `ctags` Builds an index file of function references in a C program. The `ex` and `vi` text editors can use this index file to locate the correct file for the function you name.
- `lint` Checks syntactical validity of C programs more stringently than does the C compiler.

2.1. `ctags` — Build Index File for C Functions

`ctags` builds an index file of function references in a C program. The `ex` and `vi` text editors can use this index file to locate the correct file for the function you name.

Let us look at a directory containing a program that assist in generating an index for manuals:

```
tutorial% ls index.assist
Makefile  build.index.c  index.assist.h  print.index.c
SCCS      index.assist.c  index.token.c
tutorial%
```

Now if we look inside the `Makefile` for the rule that builds the tags, we see these relevant fragments:

```

.
.
lines of Makefile
.
.
SOURCES = index.assist.c build.index.c print.index.c \
          index.token.c
.
.
more lines of Makefile
.
.
tags: $(SOURCES)
      ctags $(SOURCES)
.
.
more lines of Makefile
.
.

```

Now we run a `make tags` in that directory and we see the results:

```

tutorial% make tags
ctags index.assist.c build.index.c print.index.c \
      index.token.c
tutorial%

```

Now there is a `tags` file that acts as the index for the program. How do you use this? Suppose you want to edit the `print_index` function. You can simply say:

```

tutorial% vi -t print_index

```

The `-t` option instructs `vi` to use the `tags` file and look for the `print_index` function — then `vi` finds that the required function is in the file called `print.index.c`.

The other use of this is when you are *already editing* some file and want to look at a function that's in another file. You then use the `:ta` command of `ex`. For example, suppose you are editing the `main` function and you want to look for the `insert_index_entry` function which is in another file. You use the `:ta` command like:

```

:ta insert_index_entry

```

command and then `ex/vi` does an effective `:e` command to read in the file containing the specified function. The `insert_index_entry` function happens to be in the file called `build.index.c` and `ex/vi` announces this fact at the bottom of the screen when it reads in the appropriate file.

2.2. Lint — A C Program Checker

`lint` examines C source programs, detecting a number of bugs and obscurities. `lint` enforces the type rules of C more strictly than the C compiler. `lint` may also be used to enforce a number of portability restrictions involved in moving programs between different machines and/or operating systems. Another option detects a number of wasteful, or error-prone, constructions which nevertheless are, strictly speaking, legal.

`lint` accepts multiple input files and library specifications, and checks them for consistency.

The separation of function between `lint` and the C compilers has both historical and practical rationale. The compilers turn C programs into executable files rapidly and efficiently. This is possible in part because the compilers do not do sophisticated type checking, especially between separately compiled programs. `lint` takes a more global, leisurely view of the program, looking much more carefully at the compatibilities.

This document discusses the use of `lint`, gives an overview of its implementation, and gives some hints on writing machine-independent C code.

Using Lint

Suppose there are two C[1] source files, *file1.c* and *file2.c*, which are ordinarily compiled and loaded together. The command:

```
tutorial% lint file1.c file2.c
```

produces messages describing inconsistencies and inefficiencies in the programs. `lint` enforces the typing rules of C more strictly than the C compiler (for both historical and practical reasons) enforces them. The command:

```
tutorial% lint -p file1.c file2.c
```

produces, in addition to the types of messages described above, additional messages relating to portability of the programs to other operating systems and machines. Replacing the `-p` by `-h` produces messages about various error-prone or wasteful constructions which, strictly speaking, are not bugs. Saying `-hp` gets the whole works.

The next several sections describe the major messages; the document closes with sections discussing the implementation and giving suggestions for writing portable C. There is a summary of `lint` options in section Current Lint Options.

A Word About Philosophy

Many of the facts which `lint` needs may be impossible to discover. For example, whether a given function in a program ever gets called may depend on the input data. Deciding whether `exit` is ever called is equivalent to solving the famous ‘halting problem,’ which is known to be recursively undecidable.

Thus, most of the `lint` algorithms are a compromise. If a function is never mentioned, it can never be called. If a function is mentioned, `lint` assumes it can be called; this is not necessarily so, but in practice is quite reasonable.

`lint` tries to give information with a high degree of relevance. Messages of the form 'xxx might be a bug' are easy to generate, but are acceptable only in proportion to the fraction of real bugs they uncover. If this fraction of real bugs is too small, the messages lose their credibility and serve merely to clutter up the output, obscuring the more important messages.

Keeping these issues in mind, we now consider in more detail the classes of messages which `lint` produces.

Unused Variables and Functions

As programs evolve and develop, previously used variables and arguments to functions may become unused; it is not uncommon for external variables, or even entire functions, to become unnecessary, and yet not be removed from the source. These 'errors of commission' rarely make working programs fail, but they are a source of inefficiency, and make programs harder to understand and change. Moreover, information about such unused variables and functions can occasionally serve to discover bugs; if a function does a necessary job, and is never called, something is wrong!

`lint` complains about variables and functions which are defined but not otherwise mentioned. An exception is variables which are declared through explicit `extern` statements but are never referenced; thus the statement:

```
extern float sin();
```

will evoke no comment if `sin` is never used. Note that this agrees with the semantics of the C compiler. In some cases, these unused external declarations might be of some interest; they can be discovered by adding the `-x` option to the `lint` invocation.

Certain styles of programming require many functions to be written with similar interfaces; frequently, some of the arguments may be unused in many of the calls. The `-v` option is available to suppress the printing of complaints about unused arguments. When `-v` is in effect, no messages are produced about unused arguments except for those arguments which are unused and also declared as register arguments; this can be considered an active (and preventable) waste of the register resources of the machine.

There is one case where information about unused, or undefined, variables is more distracting than helpful. This is when `lint` is applied to some, but not all, files out of a collection which are to be loaded together. In this case, many of the functions and variables defined may not be used, and, conversely, many functions and variables defined elsewhere may be used. The `-u` option may be used to suppress the spurious messages which might otherwise appear.

Set/Used Information

`lint` attempts to detect cases where a variable is used before it is set. This is very difficult to do well; many algorithms take a good deal of time and space, and still produce messages about perfectly valid programs. `lint` detects local variables (automatic and register storage classes) whose first use appears physically earlier in the input file than the first assignment to the variable. It assumes that taking the address of a variable constitutes a 'use,' since the actual use may occur at any later time, in a data-dependent fashion.

The restriction to the physical appearance of variables in the file makes the algorithm very simple and quick to implement, since the true flow of control need not be discovered. It does mean that `lint` can complain about some programs which are legal, but these programs would probably be considered bad on stylistic grounds (for example, might contain at least two `goto`'s). Because static and external variables are initialized to 0, no meaningful information can be discovered about their uses. The algorithm deals correctly, however, with initialized automatic variables, and variables which are used in the expression which first sets them.

The set/used information also permits recognition of those local variables which are set and never used; these form a frequent source of inefficiencies, and may also be symptomatic of bugs.

Flow of Control

`lint` attempts to detect unreachable portions of the programs which it processes. It complains about unlabeled statements immediately following `goto`, `break`, `continue`, or `return` statements. An attempt is made to detect loops which can never be left at the bottom, detecting the special cases `while(1)` and `for(;;)` as infinite loops. `lint` also complains about loops which cannot be entered at the top; some valid programs may have such loops, but at best they are bad style, at worst bugs.

`lint` has an important area of blindness in the flow of control algorithm: it has no way of detecting functions which are called and never return. Thus, a call to `exit` may cause unreachable code which `lint` does not detect; the most serious effects of this are in the determination of returned function values (see the next section).

One form of unreachable statement that `lint` does not complain about is a `break` statement that cannot be reached — programs generated by `yacc`[2], and especially `lex`[3], may have literally hundreds of unreachable `break` statements. The `-O` option in the C compiler often eliminates the resulting object code inefficiency. Thus, these unreached statements are of little importance — there is typically nothing the user can do about them, and the resulting messages would clutter up the `lint` output. If these messages are desired, `lint` can be invoked with the `-b` option.

Function Values

Sometimes functions return values which are never used; sometimes programs incorrectly use function 'values' which are never returned. `lint` addresses this problem in a number of ways.

Locally, within a function definition, the appearance of both:

```
return( expr );
```

and:

```
return;
```

statements results in the message

```
function name contains return(e) and return
```

The most serious difficulty with this is detecting when a function return is

implied by flow of control reaching the end of the function. This can be seen with a simple example:

```
f ( a ) {  
    if ( a )  
        return ( 3 );  
    g ();  
}
```

Notice that, if *a* tests false, *f* will call *g* and then return with no defined return value; this will trigger a complaint from `lint`. If *g*, like `exit`, never returns, the message will still be produced when in fact nothing is wrong.

In practice, some potentially serious bugs have been discovered by this feature; it also accounts for a substantial fraction of the 'noise' messages produced by `lint`.

On a global scale, `lint` detects cases where a function returns a value, but this value is sometimes, or always, unused. When the value is always unused, it may constitute an inefficiency in the function definition. When the value is sometimes unused, it may represent bad style (for example, not testing for error conditions).

The dual problem, using a function value when the function does not return one, is also detected. This is a serious problem. Amazingly, this bug has been observed on a couple of occasions in 'working' programs; the desired function value just happened to have been computed in the function return register!

Type Checking

`lint` enforces the type checking rules of C more strictly than the compiler does. The additional checking is in four major areas: across certain binary operators and implied assignments, at the structure selection operators, between the definition and uses of functions, and in the use of enumerations.

There are a number of operators which have an implied balancing between types of the operands. The assignment, conditional (`?:`), and relational operators have this property; the argument of a `return` statement, and expressions used in initialization also suffer similar conversions. In these operations, `char`, `short`, `int`, `long`, `unsigned`, `float`, and `double` types may be freely intermixed. The types of pointers must agree exactly, except that arrays of *x*'s can, of course, be intermixed with pointers to *x*'s.

The type checking rules also require that, in structure references, the left operand of the `—>` be a pointer to structure, the left operand of the `.` be a structure, and the right operand of these operators be a member of the structure implied by the left operand. Similar checking is done for references to unions.

Strict rules apply to function argument and return value matching. The types `float` and `double` may be freely matched, as may the types `char`, `short`, `int`, and `unsigned`. Also, pointers can be matched with the associated arrays. Aside from this, all actual arguments must agree in type with their declared counterparts.

With enumerations, checks are made that enumeration variables or members are not mixed with other types, or other enumerations, and that the only operations

applied are =, initialization, ==, !=, and function arguments and return values.

Type Casts

The type casting feature in C was introduced largely as an aid to producing more portable programs. Consider the assignment:

```
p = 1 ;
```

where *p* is a character pointer. `lint` will quite rightly complain. Now, consider the assignment

```
p = (char *)1 ;
```

in which a cast has been used to convert the integer to a character pointer. The programmer obviously had a strong motivation for doing this, and has clearly signaled his intentions. It seems harsh for `lint` to continue to complain about this. On the other hand, if this code is moved to another machine, such code should be looked at carefully. The `-c` option controls the printing of comments about casts. When `-c` is in effect, casts are treated as though they were assignments subject to complaint; otherwise, all legal casts are passed without comment, no matter how strange the type mixing seems to be.

Nonportable Character Use

On the PDP-11, characters are signed quantities, with a range from -128 to 127. In most other C implementations, characters take on only positive values. Thus, `lint` will mark certain comparisons and assignments as being illegal or nonportable. For example, the fragment:

```
char c;
...
if( (c = getchar()) < 0 ) ...
```

works on the PDP-11, but will fail on machines where characters always take on positive values. The real solution is to declare *c* an integer, since `getchar` is actually returning integer values. In any case, `lint` will say 'nonportable character comparison'.

A similar issue arises with bitfields; when assignments of constant values are made to bitfields, the field may be too small to hold the value. This is especially true because on some machines bitfields are considered as signed quantities. While it may seem unintuitive to consider that a two-bit field declared of type `int` cannot hold the value 3, the problem disappears if the bitfield is declared to have type `unsigned`.

Assignments of longs to ints

Bugs may arise from the assignment of a `long` to an `int`, which may lose accuracy. This may happen in programs which have been incompletely converted to use `typedefs`. When a `typedef` variable is changed from `int` to `long`, the program can stop working because some intermediate results may be assigned to `int`'s, losing accuracy. Since there are a number of legitimate reasons for assigning `longs` to `ints`, the detection of these assignments is enabled by the `-a` option.

Strange Constructions

`lint` flags several perfectly legal, but somewhat strange, constructions — it is hoped that the messages encourage better code quality, clearer style, and may even point out bugs. The `-h` option is used to enable these checks. For example, in the statement:

```
*p++ ;
```

the `*` does nothing; this provokes the message ‘null effect’ from `lint`. The program fragment:

```
unsigned x ; if( x < 0 ) ...
```

is clearly somewhat strange; the test will never succeed. Similarly, the test:

```
if( x > 0 ) ...
```

is equivalent to:

```
if( x != 0 )
```

which may not be the intended action. `lint` will say ‘degenerate unsigned comparison’ in these cases. If one says:

```
if( 1 != 0 ) ...
```

`lint` reports ‘constant in conditional context’, since the comparison of 1 with 0 gives a constant result.

Another construction detected by `lint` involves operator precedence. Bugs which arise from misunderstandings about the precedence of operators can be accentuated by spacing and formatting, making such bugs extremely hard to find. For example, the statements:

```
if( x&077 == 0 ) ...
```

or

```
x<<2 + 40
```

probably do not do what was intended. The best solution is to parenthesize such expressions, and `lint` encourages this by an appropriate message.

Finally, when the `-h` option is in force `lint` complains about variables which are redeclared in inner blocks in a way that conflicts with their use in outer blocks. This is legal, but is considered by many (including the author) to be bad style, usually unnecessary, and frequently a bug.

Ancient History

There are several forms of older syntax which are being officially discouraged. These fall into two classes, assignment operators and initialization.

The older forms of assignment operators (for example, `+=`, `-=`, ...) could result in ambiguous expressions, such as:

```
a -=1 ;
```

which could be taken as either:

```
a -= 1 ;
```

or:

```
a = -1 ;
```

The situation is especially perplexing if this kind of ambiguity arises as the result of a macro substitution. The newer, and preferred operators (`+=`, `-=`, etc.) have no such ambiguities. To spur the abandonment of the older forms, `lint` complains about these old-fashioned operators., and the Sun C compiler issues warning messages about them.

A similar issue arises with initialization. The older language allowed:

```
int x 1 ;
```

to initialize `x` to 1, also creating syntactic difficulties. For example:

```
int x ( -1 ) ;
```

looks somewhat like the beginning of a function declaration:

```
int x ( y ) { ...
```

and the compiler must read a fair ways past `x` in order to sure what the declaration really is. Again, the problem is even more perplexing when the initializer involves a macro. The current syntax places an equals sign between the variable and the initializer:

```
int x = -1 ;
```

This is free of any possible syntactic ambiguity.

Pointer Alignment

Certain pointer assignments may be reasonable on some machines, and illegal on others, due entirely to alignment restrictions. For example, on the PDP-11, it is reasonable to assign integer pointers to double pointers, since double-precision values may begin on any integer boundary. On the Honeywell 6000, double-precision values must begin on even word boundaries; thus, not all such assignments make sense. `lint` tries to detect cases where pointers are assigned to other pointers, and such alignment problems might arise. The message ‘possible pointer alignment problem’ results from this situation whenever either the `-p` or `-h` options are in effect.

Multiple Uses and Side Effects

In complicated expressions, the best order in which to evaluate subexpressions may be highly machine-dependent. For example, on machines (like the PDP-11) in which the stack runs backwards, function arguments will probably be best evaluated from right-to-left; on machines with a stack running forward, left-to-right seems most attractive. Function calls embedded as arguments of other functions may or may not be treated similarly to ordinary arguments. Similar issues arise with other operators which have side effects, such as the assignment operators and the increment and decrement operators.

In order that the efficiency of C on a particular machine not be unduly compromised, the C language leaves the order of evaluation of complicated expressions up to the local compiler, and, in fact, the various C compilers have considerable differences in the order in which they will evaluate complicated expressions. In particular, if any variable is changed by a side effect, and also used elsewhere in the same expression, the result is explicitly undefined.

`lint` checks for the important special case where a simple scalar variable is affected. For example, the statement:

```
a[i] = b[i++] ;
```

will draw the complaint:

```
warning: i evaluation order undefined
```

Implementation

`lint` consists of two programs and a driver. The first program is a version of the Portable C Compiler[4], [5] which is the basis of many C compilers, including Sun's. This compiler does lexical and syntax analysis on the input text, constructs and maintains symbol tables, and builds trees for expressions. Instead of writing an intermediate file which is passed to a code generator, as the compilers do, `lint` produces an intermediate file which consists of lines of ASCII text. Each line contains an external variable name, an encoding of the context in which it was seen (use, definition, declaration, etc.), a type specifier, and a source file name and line number. The information about variables local to a function or file is collected by accessing the symbol table, and examining the expression trees.

Comments about local problems are produced as detected. The information about external names is collected onto an intermediate file. After all the source files and library descriptions have been collected, the intermediate file is sorted to bring all information collected about a given external name together. The second, rather small, program then reads the lines from the intermediate file and compares all of the definitions, declarations, and uses for consistency.

The driver controls this process, and is also responsible for making the options available to both passes of `lint`.

Portability

C on the Honeywell and IBM systems is used, in part, to write system code for the host operating system. This means that the implementation of C tends to follow local conventions rather than adhere strictly to UNIX system conventions.

Despite these differences, many C programs have been successfully moved to GCOS and the various IBM installations with little effort. This section describes some of the differences between the implementations, and discusses the `lint` features which encourage portability.

Uninitialized external variables are treated differently in different implementations of C. Suppose two files both contain a declaration without initialization, such as:

```
int a ;
```

outside of any function. The UNIX loader resolves these declarations, and sets aside only a single word of storage for *a*. Under the GCOS and IBM implementations, this is not feasible (for various stupid reasons!) so each such declaration sets aside a word of storage called *a*. When loading or library editing takes place, this creates fatal conflicts which prevent the proper operation of the program. `lint` detects such multiple definitions if it is invoked with the `-p` option.

A related difficulty comes from the amount of information retained about external names during the loading process. On the UNIX system, externally known names have seven significant characters, with the upper/lower case distinction kept. On the IBM systems, there are eight significant characters, but the case distinction is lost. On GCOS, there are only six characters, of a single case. This leads to situations where programs run on the UNIX system, but encounter loader problems on the IBM or GCOS systems. `lint -p` maps all external symbols to one case and truncates them to six characters, providing a worst-case analysis.

A number of differences arise in the area of character handling: characters in the UNIX system are eight bit ASCII, while they are eight bit EBCDIC on the IBM, and nine bit ASCII on GCOS. Moreover, character strings go from high to low bit positions ('left to right') on GCOS and IBM, and low to high ('right to left') on the PDP-11. This means that code attempting to construct strings out of character constants, or attempting to use characters as indices into arrays, must be looked at with great suspicion. `lint` is of little help here, except to option multi-character character constants.

Of course, the word sizes are different! This is less troublesome than might be expected, at least when moving from the UNIX system (16 bit words) to the IBM (32 bits) or GCOS (36 bits). The main problems are likely to arise in shifting or masking. C now supports a bit-field facility, which can be used to write much of this code in a reasonably portable way. Frequently, portability of such code can be enhanced by slight rearrangements in coding style. Many of the incompatibilities seem to have the flavor of writing:

```
x &= 0177700 ;
```

to clear the low order six bits of `x`. This suffices on the PDP-11, but fails badly on GCOS and IBM. If the bit field feature cannot be used, the same effect can be obtained by writing:

```
x &= ~ 077 ;
```

which will work on all these machines.

The right shift operator is arithmetic shift on the PDP-11, and logical shift on most other machines. To obtain a logical shift on all machines, the left operand can be typed `unsigned`. Characters are considered signed integers on the PDP-11, and unsigned on the other machines. This persistence of the sign bit may be reasonably considered a bug in the PDP-11 hardware which has infiltrated itself into the C language. If there were a good way to discover the programs which would be affected, C could be changed; in any case, `lint` is no help here.

The above discussion may have made the problem of portability seem bigger than it in fact is. The issues involved here are rarely subtle or mysterious, at least to the implementor of the program, although they can involve some work to straighten out. The most serious bar to the portability of UNIX system utilities has been the inability to mimic essential UNIX system functions on the other systems. The inability to seek to a random character position in a text file, or to establish a pipe between processes, has involved far more rewriting and debugging than any of the differences in C compilers. On the other hand, `Lint` has been very helpful in moving the UNIX operating system and associated utility

programs to other machines.

Shutting Lint Up

There are occasions when the programmer is smarter than `lint`. There may be valid reasons for 'illegal' type casts, functions with a variable number of arguments, etc. Moreover, as specified above, the flow of control information produced by `lint` often has blind spots, causing occasional spurious messages about perfectly reasonable programs. Thus, some way of communicating with `lint`, typically to shut it up, is desirable.

The form which this mechanism should take is not at all clear. New keywords would require current and old compilers to recognize these keywords, if only to ignore them. This has both philosophical and practical problems. New preprocessor syntax suffers from similar problems.

What was finally done was to make `lint` recognize a number of words when they were embedded in comments. This required minimal preprocessor changes; the preprocessor just had to agree to pass comments through to its output, instead of deleting them as had been previously done. Thus, `lint` directives are invisible to the compilers, and the effect on systems with the older preprocessors is merely that the `lint` directives don't work.

The first directive is concerned with flow of control information; if a particular place in the program cannot be reached, but this is not apparent to `lint`, this can be asserted by placing the directive

```
/* NOTREACHED */
```

just before that spot in the program. Similarly, if it is desired to turn off strict type checking for the next expression, the directive

```
/* NOSTRICT */
```

can be used; the situation reverts to the previous default after the next expression. The `-v` option can be turned on for one function by the directive:

```
/* ARGSUSED */
```

Complaints about variable numbers of arguments in calls to a function can be turned off by the directive:

```
/* VARARGS */
```

preceding the function definition. In some cases, it is desirable to check the first several arguments, and leave the later arguments unchecked. This can be done by following the `VARARGS` keyword immediately with a digit giving the number of arguments which should be checked; thus,

```
/* VARARGS2 */
```

checks the first two arguments and leaves the others unchecked. Finally, the directive:

```
/* LINTLIBRARY */
```

at the head of a file identifies this file as a library declaration file; this topic is worth a section by itself.

Library Declaration Files

`lint` accepts certain library directives, such as:

```
-ly
```

and tests the source files for compatibility with these libraries. This is done by accessing library description files whose names are constructed from the library directives. These files all begin with the directive:

```
/* LINTLIBRARY */
```

which is followed by a series of dummy function definitions. The critical parts of these definitions are the declaration of the function return type, whether the dummy function returns a value, and the number and types of arguments to the function. The `VARARGS` and `ARGSUSED` directives can be used to specify features of the library functions.

`lint` library files are processed almost exactly like ordinary source files. The only difference is that functions which are defined in a library file, but not used in a source file, draw no complaints. `lint` does not simulate a full library search algorithm, and complains if the source files contain a redefinition of a library routine (this is a feature!).

By default, `lint` checks the routines it is given against a standard library file, which contains descriptions of the programs which are normally loaded when a C program is run. When the `-p` option is in effect, another file is checked containing descriptions of the standard I/O library routines which are expected to be portable across various machines. The `-n` option can be used to suppress all library checking.

Bugs, etc.

`lint` was a difficult program to write, partially because it is closely connected with matters of programming style, and partially because users usually don't notice bugs which cause `lint` to miss errors which it should have caught. By contrast, if `lint` incorrectly complains about something that is correct, the programmer reports that immediately!

A number of areas remain to be further developed. The checking of structures and arrays is rather inadequate; size incompatibilities go unchecked, and no attempt is made to match up structure and union declarations across files. Some stricter checking of the use of `typedef` is clearly desirable, but what checking is appropriate, and how to carry it out, is still to be determined.

`lint` shares the preprocessor with the C compiler. At some point it may be appropriate for a special version of the preprocessor to be constructed which checks for things such as unused macro definitions, macro arguments which have side effects which are not expanded at all, or are expanded more than once, etc.

The central problem with `lint` is the packaging of the information which it collects. There are many options which serve only to turn off, or slightly modify, certain features. There are pressures to add even more of these options.

In conclusion, it appears that the general notion of having two programs is a good one. The compiler concentrates on quickly and accurately turning the program text into bits which can be run; `lint` concentrates on issues of portability, style, and efficiency. `lint` can afford to be wrong, since incorrectness and over-

conservatism are merely annoying, not fatal. The compiler can be fast since it knows that `lint` will cover its flanks. Finally, the programmer can concentrate at one stage of the programming process solely on the algorithms, data structures, and correctness of the program, and then later retrofit, with the aid of `lint`, the desirable properties of universality and portability.

Current Lint Options

The `lint` command currently has the form

```
tutorial% lint [-abchnpsuvx ] filename . . . library-descriptors . . .
```

The options are

- a Report assignments of `long` to `int` or shorter
- b Report unreachable `break` statements
- c Complain about questionable casts
- h Perform heuristic checks
- n Do not do library checking
- p Perform portability checks
- s Same as `h` (for historical reasons)
- u Don't report unused or undefined externals
- v Don't report unused arguments
- x Report unused external declarations

Make — Maintaining Computer Programs

Make — Maintaining Computer Programs	49
3.1. Basic Features	50
Default Target	51
3.2. Description Files	52
Comments in <code>makefile</code>	52
Continuation Lines	52
Include Lines	53
Macro Definitions	53
Using Macros	53
Translations in Macro References	53
Recursive Makefiles	54
Entries — Dependency Lines and Rules	54
Dynamic Dependency Parameters	55
Implicit Macros	56
3.3. Using the <code>make</code> Command	56
Assigning Macros and Variables	57
Options for the <code>make</code> Command" 58 3 3	58
3.4. Implicit Rules	59
SCCS File Names	60
Built In Names and Options	61
3.5. Example	61
3.6. Suggestions and Warnings	63
3.7. Making Archive Libraries	64

3.8. Suffixes and Transformation Rules	65
Null Suffix	65
Standard Suffix List	65

Make — Maintaining Computer Programs

It is common practice to divide large programs into smaller, more manageable pieces. The pieces may require quite different treatments: some may need to be run through a macro processor, and some may need to be processed by a sophisticated program generator (for example, Yacc[1] or Lex[2]). The outputs of these generators may have to be compiled with special options and with certain definitions and declarations. The code resulting from these transformations may then need to be loaded together with certain libraries under the control of special options. Related maintenance activities involve running complicated test scripts and installing validated modules. Unfortunately, it is very easy for a programmer to forget which files depend on which others, which files have been modified recently, and the exact sequence of operations needed to make or exercise a new version of the program. After a long editing session, one may easily lose track of which files have been changed and which object modules are still valid, since a change to a declaration can obsolete a dozen other files. Forgetting to compile a routine that has been changed or that uses changed declarations usually results in a program that will not work, and a bug that can be very hard to track down. On the other hand, recompiling everything in sight just to be safe is very wasteful.

`make` mechanizes many of the activities of program development and maintenance. `make` provides a simple mechanism for maintaining up-to-date versions of programs that result from many operations on a number of files. It is possible to tell `make` the sequence of commands that create certain files, and the list of files that require other files to be current before the operations can be done. Whenever a change is made in any part of the program, `make` will create the proper files simply, correctly, and with a minimum amount of effort.

Basic Ideas

The basic operation of `make` is to find the name of a needed target in the description, ensure that all of the files on which it depends exist and are up-to-date, and then create the target if it has not been modified since its generators were. The description file really defines the graph of dependencies; `make` does a depth-first search of this graph to determine what work is really necessary.

`make` also provides a simple macro substitution facility and the ability to encapsulate commands in a single file for convenient administration.

If the information on inter-file dependences and command sequences is stored in a file, the simple command:

```
tutorial% make
```

is frequently sufficient to update the interesting files, regardless of the number that have been edited since the last 'make'. In most cases, the description file is easy to write and changes infrequently. It is usually easier to type the `make` command than to issue even one of the needed operations, so the typical cycle of program development operations becomes

think — edit — make — test ...

`make` is most useful for medium-sized programming projects; it does not solve the problems of maintaining multiple source versions⁴ or of describing huge programs.

3.1. Basic Features

The basic operation of `make` is to update a target file by ensuring that all of the files on which it depends exist and are up to date, then creating the target if it has not been modified since its dependents were. `make` does a depth-first search of the graph of dependences. The operation of the command depends on the ability to find the date and time that a file was last modified.

To illustrate, let us consider a simple example: A program named `prog` is made by compiling and loading three C-language files `x.c`, `y.c`, and `z.c` with the `lm` library. By convention, output of the C compilations is found in files named `x.o`, `y.o`, and `z.o`. Assume that the files `x.c` and `y.c` share some declarations in a file named `defs`, but that `z.c` does not. That is, `x.c` and `y.c` have the line

```
#include "defs"
```

The following text describes the relationships and operations:

```
prog: x.o y.o z.o
    cc x.o y.o z.o -lm -o prog

x.o y.o: defs
```

If this information were stored in a file named `makefile`, the command:

```
tutorial% make
```

would perform the operations needed to recreate `prog` after any changes had been made to any of the four source files `x.c`, `y.c`, `z.c`, or `defs`.

`make` operates using three sources of information: a user-supplied description file (as above), filenames and 'last-modified' times from the file system, and built-in rules to bridge some of the gaps. In our example, the first line says that `prog` depends on three '.o' files. Once these object files are current, the second line describes how to load them to create `prog`. The third line says that `x.o`

⁴ See the description of the Source Code Control System (SCCS) later in this book, for a tool for maintaining multiple source versions.

and `y.o` depend on the file `defs`. From the file system, `make` discovers that there are three `.c` files corresponding to the needed `.o` files, and uses built-in information on how to generate an object from a source file (*that is*, issue a `cc -c` command).

The following long-winded description file is equivalent to the one above, but takes no advantage of `make`'s innate knowledge:

```
prog : x.o y.o z.o
      cc x.o y.o z.o -lm -o prog
x.o : x.c defs
      cc -c x.c
y.o : y.c defs
      cc -c y.c
z.o : z.c
      cc -c z.c
```

If none of the source or object files had changed since the last time `prog` was made, all of the files would be current, and the command:

```
tutorial% make
```

would just announce this fact and stop. If, however, the `defs` file had been edited, `x.c` and `y.c` (but not `z.c`) would be recompiled, and then `prog` would be created from the new `.o` files. If only the file `y.c` had changed, only it would be recompiled, but it would still be necessary to reload `prog`.

Default Target

If no target name is given on the `make` command line, the first target mentioned in the description is created; otherwise the specified targets are made. The command:

```
tutorial% make x.o
```

would recompile `x.o` if `x.c` or `defs` had changed.

If the file exists after the commands are executed, its time of last modification is used in further decisions; otherwise the current time is used. It is often quite useful to include rules with mnemonic names and commands that do not actually produce a file with that name. These entries can take advantage of `make`'s ability to generate files and substitute macros. Thus, an entry `'save'` might be included to copy a certain set of files, or an entry `'cleanup'` might be used to throw away unneeded intermediate files. In other cases one may maintain a zero-length file purely to keep track of the time at which certain actions were performed. This technique is useful for maintaining remote archives and listings.

`make` has a simple macro mechanism for substituting in dependency lines and command strings. Macros are defined by command arguments or description file lines with embedded equal signs. A macro is invoked by preceding the name by a dollar sign; macro names longer than one character must be parenthesized. The name of the macro is either the single character after the dollar sign or a name inside parentheses. The following are valid macro invocations:

```
$(CFLAGS) $2 $(xy) $Z $(Z)
```

The last two invocations are identical. `$$` produces a dollar sign. All of these macros are assigned values during input, as shown below. Four special macros change values during the execution of the command: `$$`, `$$@`, `$$?`, and `$$<`. They are discussed below. The following fragment shows how macros are used:

```
OBJECTS = x.o y.o z.o
LIBES = -lm
prog: $(OBJECTS)
    cc $(OBJECTS) $(LIBES) -o prog
. . .
```

The command:

```
tutorial% make
```

loads the three object files with the `lm` library. The command:

```
tutorial% make "LIBES= -ll -lm"
```

loads them with both the `lex` (`-ll`) and mathematical (`-lm`) libraries, since macro definitions on the command line override definitions in the description. It is necessary to quote arguments with embedded blanks in UNIX commands.

The following sections detail the form of description files and the command line, and discuss options and built-in rules in more detail.

3.2. Description Files

A make description file, also known as a *makefile*, contains five types of information:

- Comments,
- `include` lines,
- Macro definitions,
- Dependency information,
- Executable commands.

The last two items are actually combined into a make *entry*.

Comments in *makefile*

make's comment convention is simple: all characters after a sharp (`#`) to the end of the line are ignored, as is the sharp itself. Blank lines and lines beginning with a sharp are totally ignored.

Continuation Lines

If a non-comment line is too long, it can be continued using a backslash. If the last character of a line is a backslash, the backslash, newline, and following blanks and tabs are replaced by a single blank.

Include Lines

make supports a facility for including other files into the body of a `makefile`. If the string `include` appears as the first seven letters of a line in a `makefile` and is followed by a space or a tab, the string following the word `include` is taken as a filename which the current invocation of `make` will read. `include` files can be nested to a depth of no more than about 16.

Macro Definitions

make supplies a simple macro capability. A macro *definition* is a line containing an equal sign not preceded by a colon or a tab. The name (string of letters and digits) to the left of the equal sign (trailing blanks and tabs are stripped) is assigned the string of characters following the equal sign (leading blanks and tabs are stripped, but trailing ones are not). The following are valid macro definitions:

```
A = xyz
LIBS = -lcore -lpixrect
OFFSET =
```

The last definition assigns `OFFSET` the null string. A macro that is never explicitly defined has the null string as value.

Macro definitions may also appear on the `make` command line when you actually use the `make` command (see below).

Using Macros

If *macro_name* is the name of a `make` macro, you access the definition of that macro in the body of a `makefile` with the construct

```
$macro_name
```

if *macro_name* is only a single character. If *macro_name* is longer than one character you use either of the two alternative notations:

```
$(macro_name)
or
${macro_name}
```

Taking our macro definition examples from above, you reference the `A` macro as:

```
$A
```

to generate the string `xyz`, and you reference the `LIBS` macros with one of the two alternative forms:

```
$(LIBS)
or
${LIBS}
```

to obtain the string `-lcore -lpixrect`

Translations in Macro References

There is also a facility to perform translations when a macro is referenced and evaluated. The general syntax of such a macro reference is:

```
$(macro_name : string_1 = string_2)
```

This is interpreted as:

- The macro specified by *macro_name* is evaluated, and then:

- For each occurrence of *string_1* in the evaluated macro, substitute *string_2*.

What constitutes an occurrence of *string_1* in the evaluated macro? The evaluated macro is considered to be a set of strings each separated by whitespace (spaces or tabs). An occurrence of *string_1* in the evaluated macro means that a regular expression of this form has been found in the evaluated macro:

```
.*<string_1>[tab|space]
```

There is an example of how this is used later on.

Recursive Makefiles

`makefile`'s can be set up so that they perform recursive invocations of `make`. If the sequence `$(MAKE)` appears anywhere in a Shell command line, the line is executed even if the `-n` option was specified on the original `make` command line. The `-n` option is exported across invocations of `make` (via the `MAKEFLAGS` variable), so the only thing that gets executed is the `make` command itself. You can use this feature when a hierarchy of `makefile`'s describes a collection of subsystems. You can type `make -n` and everything that *would* happen is displayed without actually executing the commands. Because of the `$(MAKE)` sequence, the lower level `make`'s get executed.

Entries — Dependency Lines and Rules

The major piece of information in a `makefile` is an *entry*. An entry consists of a *target* and *rules*. A target contains any number of target names and optional *dependency* information. A dependency specifies a set of things that the given target depends on — that is, do something to construct the target if the things it depends on have been updated since the last time the target was constructed. The general form of an entry is:

```
target-name . . . :[:] [dependent . . .] [; commands] [# . . .]
[(tab) commands] [# . . .]
. . .
```

Items inside brackets may be omitted. Targets and dependents are strings of letters, digits, periods, and slashes. Shell metacharacters '*' and '?' are expanded.

Note that a command must be preceded by a *tab* character at the beginning of the line. This is one of `make`'s less obvious and more irritating 'features'.

A command is any string of characters not including a sharp (except in quotes) or newline. Commands may appear either after a semicolon on a dependency line or on lines *beginning with a tab* immediately following a dependency line.

`make` remembers embedded newlines and tabs in sequences of Shell commands. So if you write a `for` loop in the `makefile` with indentation, `make` retains the indentation and backslashes when the commands are displayed. The output can still be piped to the Shell and is readable.

A dependency line may have either a single or a double colon. A target name may appear on more than one dependency line, but all of those lines must be of the same (single or double colon) type.

1. For the usual single-colon case, at most one of these dependency lines may have a command sequence associated with it. If the target is out of date with any of the dependents on any of the lines, and a command sequence is specified (even a null one following a semicolon or tab), it is executed;

otherwise a default creation rule may be invoked.

2. In the double-colon case, a command sequence may be associated with each dependency line; if the target is out of date with any of the files on a particular line, the associated commands are executed. A built-in rule may also be executed. This detailed form is of particular value in updating archive-type files.

If a target must be created, the sequence of commands is executed. Normally, each command line is displayed and then passed to a separate invocation of the Shell after substituting for macros. The displaying is suppressed in silent mode or if the command line begins with an @ sign. `make` normally stops if any command signals an error by returning a non zero error code.

`make` ignores errors if the `-i` option has been specified on the `make` command line, if the fake target name `.IGNORE` appears in the description file, or if the command string in the description file begins with a hyphen — these criteria are necessary because some UNIX commands return meaningless status.

Because each command line is passed to a separate invocation of the Shell, care must be taken with certain commands (for example, `cd` and Shell control commands) that have meaning only within a single Shell process; the results are forgotten before the next line is executed.

Dynamic Dependency Parameters

The dynamic dependency parameter is referenced by the `$$@` notation. This dynamic dependency parameter only has meaning on the dependency line in a `makefile`. The `$$@` refers to the current ‘thing’ to the left of the colon — the ‘thing’ to the left of the colon is the `$$` implicit macro defined below. You can also use the form `$$ (@F)` which refers to the file part of `$$@`.

How do you use this form? Well suppose you have a program called `buzz`. You can refer to `buzz` in your `makefile` like this:

```
buzz: $$@.c
```

This means that `buzz` depends on `buzz.c`. This dynamic dependency parameter finds most use in maintaining a bunch of programs that only depend on a single source file. Suppose you have a directory with many small toy programs. You could have a `makefile` that looks something like this:

```
PROGRAMS = buzz biorythm checkbook tictactoe

$(PROGRAMS): $$@.c
$(CC) -o $$? -o $$@
```

The second form of the dynamic dependency parameter using the `$$ (@F)` notation finds most use when maintaining some directory from the contents of another directory. Suppose the source files of `/usr/include` reside in `/usr/src/usr.include`. What you want is that every time you update one of the `.h` file in `/usr/src/usr.include`, then type `make`, the appropriate file gets moved into the `/usr/include` directory. Here is a fragment of a `makefile` (residing in the `/usr/src/usr.include` directory) that would do this job:

```

DESTDIR = /usr/include           define the destination (target) directory

SOURCE_FILES = $(DESTDIR)/a.out.h \ define all the dependents
$(DESTDIR)/ar.h \
$(DESTDIR)/assert.h \
$(DESTDIR)/cgicbind.h \
$(DESTDIR)/cgiconstants.h \
.
.
.
$(DESTDIR)/time.h \
$(DESTDIR)/usercore.h \
$(DESTDIR)/utmp.h \
$(DESTDIR)/varargs.h \
$(DESTDIR)/vfont.h

$(SOURCE_FILES): $$(@F)           now here is the target and the rule
cp $? $@
chmod 0444 $@

```

Implicit Macros

`make` reads the user environment (see section 3.3 below), and sets certain macros before issuing any command:

`$@` is set to the name of the file to be 'made'.

`$?` is set to the string of names that were found to be younger than the target. If the command was generated by an implicit rule (see below),

`$<` is the name of the related file that caused the action, and

`$*` is the prefix shared by the current and the dependent filenames.

These implicit macros are useful generic terms for current targets and out-of-date relatives. There are some extra forms of the macros, namely, `$(@D)`, `$(@F)`, `$(*D)`, `$(*F)`, `$(<D)`, and `$(<F)`. For each of the macros, the `D` part refers to the *Directory* part of the name, and the `F` part refers to the *File* part of the name. These macros are used when building hierarchical `makefile`'s. They provide access to directory names for using the `cd` command of the Shell. For example, a Shell command could be:

```
cd $(<D); $(MAKE) $(<F)
```

If a file must be made but there are no explicit commands or relevant built-in rules, the commands associated with the name `.DEFAULT` are used. If there is no such name, `make` displays a message and stops.

3.3. Using the `make` Command

The `make` command takes three kinds of arguments: macro definitions, options, and target filenames. In addition, `make` obtains information from the environment.

```
tutorial% make [ options ] [ macro definitions ] [ targets ]
```

The following summary of the operation of the command explains how these arguments are interpreted.

Assigning Macros and Variables

`make` reads environment variables and adds them to the macro definitions every time the command executes. `make` maintains a macro called `MAKEFLAGS`, which is a string defined as the collection of all command line options (*sans* their minus signs). The `MAKEFLAGS` macro is exported and is thus accessible to further invocations of `make`. Here is how `make` assigns macro definitions:

1. Read the `MAKEFLAGS` environment variable. If `MAKEFLAGS` does not exist or is null, set `MAKEFLAGS` to the null string. Otherwise, each letter in `MAKEFLAGS` is taken to be a command line option and is processed as such. The `-f`, `-p`, and `-r` options do not get processed.
2. Read options from the command line. Options from the command line add to the previous settings from the `MAKEFLAGS` environment variable.
3. Read macro definitions from the command line. Such macro definitions are made *non-resettable* and any further assignments to these names are ignored.
4. Read `make`'s internal list of macro definitions. Table 3-3 shows the built-in macro names and their defaults.
5. Read the environment. Environment variables are treated as macro definitions and are *exported*. Now because `MAKEFLAGS` is not a `make` internal variable, this has the effect of doing the same assignment twice. The exception to this is when `MAKEFLAGS` is assigned on the command line. The reason for reading `MAKEFLAGS` first is to turn on the debug option (if the debug option was indeed specified) before doing anything else.
6. Read the `makefile`'s. Assignments in the `makefile`'s override the environment, unless you used the `-e` command line option to tell `make` to have the environment override assignments made in the `makefile`'s.

Here is a summary of how the various parts of the environment, internal definitions, command line options, and the contents of `makefile`'s are assigned. The order of assignment is from the least binding to the most binding — that is, higher numbered items override lower numbered items.

Table 3-1 *Summary of Assigning Macros and Variables*

<i>-e option not specified</i>		<i>-e option specified</i>	
1	internal definitions	1	internal definitions
2	environment	2	<code>makefile(s)</code>
3	<code>makefile(s)</code>	3	environment
4	command line	4	command line

Options for the `make` Command"

Next, the options are examined. The permissible options are:

`-f filename`

Use *filename* as the name of the description file. A file name of `-` denotes the standard input. In the absence of the `-f` option, `make` looks for a set of standard filenames as follows:

- `makefile` in the current directory,
- `Makefile` in the current directory.
- `s.makefile` in the current directory.
- `s.Makefile` in the current directory.
- `SCCS/s.makefile`,
- `SCCS/s.Makefile`.

The contents of description files specified by the `-f` option override the built-in rules if they are present.

- `-p` Print out the complete set of macro definitions and target descriptions.
- `-i` Ignore error codes returned by invoked commands. This mode is entered if the fake target name `.IGNORE` appears in the description file.
- `-k` Abandon work on the current entry, but continue on other branches that do not depend on that entry.
- `-s` Silent mode. Do not print command lines before executing. This mode is also entered if the fake target name `.SILENT` appears in the description file.
- `-r` Do not use the built-in rules.
- `-n` No execute mode. Print commands, but do not execute them. Even lines beginning with an `@` are printed.
- `-b` Compatibility mode for old makefiles. `-b` is on by default.
- `-e` Environment variables override assignments within makefiles.
- `-t` Touch the target files (bringing them up-to-date) rather than issue the usual commands.
- `-d` Debug mode. Print out detailed information on files and times examined.
- `-q` Question. The `make` command returns a zero or non-zero status code depending on whether the target file is or is not up-to-date.
- `-S` Undoes the effect of the `-k` option.

Remaining arguments are assumed to be the names of targets to be made; they are done in left-to-right order. If there are no such arguments, the first name in the description files that does not begin with a period is 'made'.

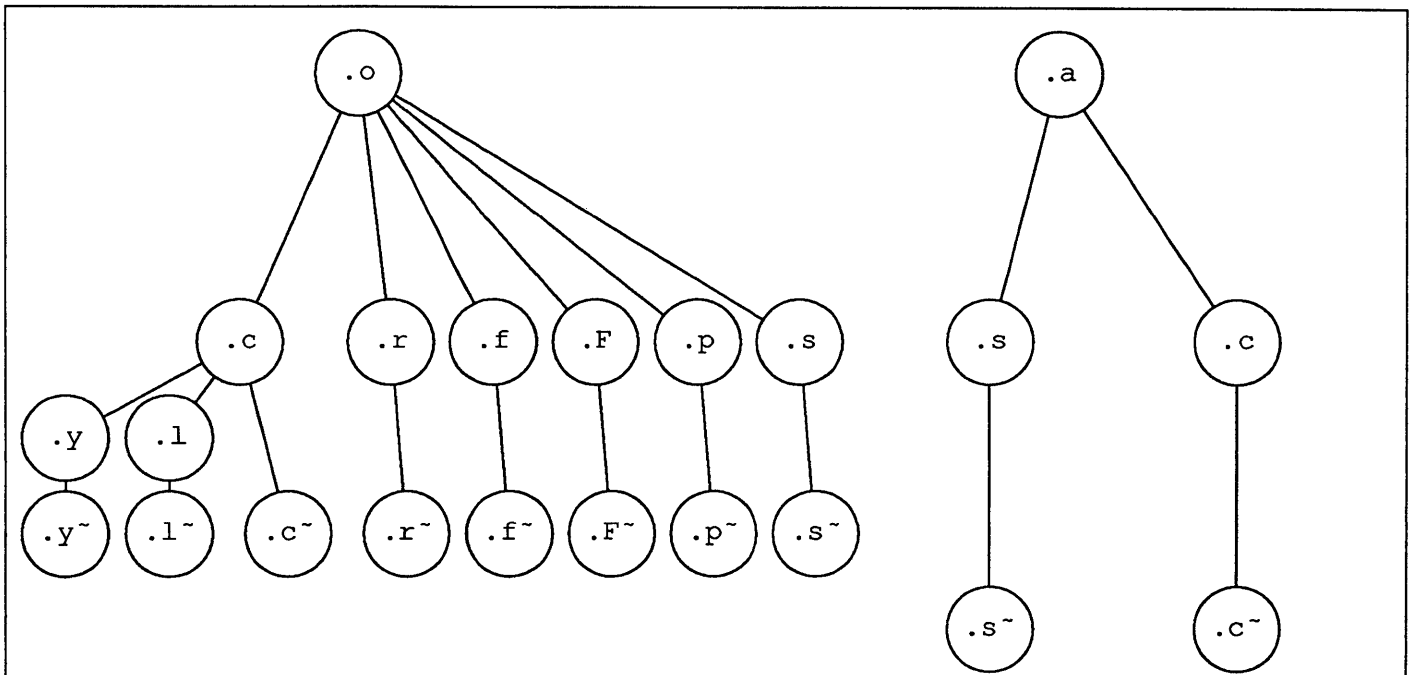
3.4. Implicit Rules

The `make` program uses a table of interesting suffixes and a set of transformation rules to supply default dependency information and implied commands. Section 3.8 describes these tables and means of overriding them. The default suffix list is:

Table 3-2 *Default Suffix List for Make*

<i>Suffix</i>	<i>Type of File</i>
.o	Object file
.c	C source file
.c~	C source file from SCCS s-file
.r	Ratfor source file
.r~	Ratfor source file from SCCS s-file
.f	Fortran source file
.f~	Fortran source file from SCCS s-file
.F	Fortran source file
.F~	Fortran source file from SCCS s-file
.s	Assembler source file
.s~	Assembler source file from SCCS s-file
.y	Yacc-C source grammar
.y~	Yacc-C source file from SCCS s-file
.p	Pascal source
.p~	Pascal source file from SCCS s-file
.l	Lex source grammar
.l~	Lex source grammar from SCCS s-file
.h	Include file
.h~	Include file from SCCS s-file
.sh	Shell script
.sh~	Shell script from SCCS s-file

The following diagram summarizes the default transformation paths. If there are two paths connecting a pair of suffixes, the longer one is used only if the intermediate file exists or is named in the description.



To get from a given type of source file to the `.o` file the appropriate compiler is called up to generate the `.o` file. If there is an SCCS version of the source file available, the `sccs get` command is called first, followed by the appropriate compiler. Notice that there are also transformation rules to create a library (`.a` files) from source as well.

If the file `x.o` were needed and there were an `x.c` in the description or directory, it would be compiled. If there were also an `x.l`, that grammar would be run through `lex` before compiling the result. However, if there were no `x.c` but there were an `x.l`, `make` would discard the intermediate C-language file and use the direct link in the graph above.

SCCS File Names

The syntax of `make` doesn't permit referencing filenames that have prefixes directly. This is all right for most UNIX system filenames since most reasonable people use suffixes to distinguish different kinds of files — `.c` for C source files, `.f` for FORTRAN source files, and so on. SCCS database files are a glaring exception to the conventions — SCCS database filenames are *refixed* with `.s`. To avoid redefining the syntax for naming rules, `make` employs a trick — the tilde character (`~`) is used to identify SCCS database files. Thus, `.c~.o` refers to the rule for making a `.o` file out of a C language source file that's stored in an SCCS `.s` file. Specifically, the rule in this case is:

```
.c~.o:
$(GET) -G$*.c $(GFLAGS) $<
$(CC) $(CFLAGS) -c $*.c
```

So, a tilde appended to any suffix transforms the file search into an SCCS file name search with the actual suffix named by the dot and all characters up to (but not including) the tilde.

Built In Names and Options

It is possible to change the names of some of the compilers used in the default, or the option arguments with which they are invoked by knowing the macro names used. The compiler names and the options passed to them are determined by the macros as in the table below:

Table 3-3 *Built In Compiler Names and Options*

<i>Macro Name</i>	<i>Default Value</i>	<i>Description</i>
MAKE	make	Name of the make command
YACC	yacc	Name of the yacc command
YFLAGS	<i>null</i>	options for the yacc command
LEX	lex	Name of the lex command
LFLAGS	<i>null</i>	options for the lex command
LD	ld	Name of the link editor
LDFLAGS	<i>null</i>	options for the link editor
CC	cc	Name of the C compiler
CFLAGS	<i>null</i>	options for the C compiler
FC	f77	Name of the FORTRAN 77 compiler
FFLAGS	<i>null</i>	options for the FORTRAN 77 compiler
AS	as	Name of the Assembler
ASFLAGS	<i>null</i>	options for the Assembler
GET	/usr/sccs/get	Name of the sccs get command
GFLAGS	<i>null</i>	options for the sccs get command

The command:

```
tutorial% make CC=newcc
```

uses the `newcc` command instead of the usual C compiler. The macros `CFLAGS`, `FFLAGS`, `PFLAGS`, `RFLAGS`, `YFLAGS`, and `LFLAGS` may be set to issue these commands with optional options.

```
tutorial% make "CFLAGS=-O"
```

uses the optimizing C compiler.

The `make` variable `MFLAGS` is also useful — it contains a list of the command-line arguments given to this invocation of `make`.

3.5. Example

As an example of the use of `make`, consider the following description file which could be used to maintain the `make` command itself. The code for `make` is spread over a number of C source files and a yacc grammar. The description file contains:

```
#
```

```
# @(#)Makefile 1.5 85/07/08 SMI; from S5R2 1.7
#

# The rules.c file can be modified locally for
# people who still like things like fortran.

LDFLAGS =
INSDIR = $(DESTDIR)/bin
LIBS =
CFLAGS = -O -DBSD -DSCCSDIR

OBJECTS = \
    main.o \
    doname.o \
    misc.o \
    files.o \
    rules.o \
    dosys.o \
    gram.o \
    dyndep.o \
    prtmem.o

all: make

make: $(OBJECTS)
    $(CC) -o make $(LDFLAGS) $(OBJECTS) $(LIBS)

gram.c: gram.y

gram.o: gram.c

$(OBJECTS): defs

install: all
    install -c -s make $(INSDIR)

clean:
    -rm -f *.o a.out core errs make gram.c

tags: NOW
    ctags *.ch

NOW:
    ;

$(GET) $(GFLAGS) -p s.$< > $<
```

Although none of the source files or grammars were mentioned by name in the description file, make finds them using its suffix rules and issued the needed commands.

3.6. Suggestions and Warnings

The most common difficulties arise from `make`'s specific meaning of dependency. If file `x.c` has a `#include "defs"` line, the object file `x.o` depends on `defs`; the source file `x.c` does not. If `defs` is changed, it is not necessary to do anything to the file `x.c`, while it is necessary to recreate `x.o`.

`-n` (no execute) Option

To discover what `make` would do, the `-n` option is very useful. The command:

```
tutorial% make -n
```

orders `make` to display the commands it would issue without actually executing them. See section 3.2.7 earlier for other ramifications of using the `-n` options.

`-t` (touch) Option

If a change to a file is absolutely certain to be benign (for example, adding a new definition to an include file), the `-t` (touch) option can save a lot of time: instead of issuing a large number of superfluous recompilations, `make` updates the modification times on the affected file. Thus, the command:

```
tutorial% make -ts
```

(‘touch silently’) makes the relevant files appear up-to-date. Obvious care is necessary, since this mode of operation subverts the intention of `make` and destroys all memory of the previous relationships.

`-d` (debug) Option

The debugging option (`-d`) generates a very detailed description of what `make` is doing, including the file times. The output is verbose, and recommended only as a last resort.

Compiler and Loader Options

Another common blunder is specifying some option for the compiler but forgetting it on the linker. You might have this fragment in a `makefile`:

```

:
:
  lines of makefile
:
:
CFLAGS = -g           to get the debug option for dbx
:
:
  lines of makefile
:
:
prog: s.o t.o
  cc -o prog s.o t.o
:
:
  lines of makefile
:
:

```

and think that this will work. It won't because `CFLAGS` only applies to the `cc`

`-c s.c` part of the compilation and not the `cc -o prog s.o t.o` of the compile. And `dbx` won't work unless you specified the `-g` option for both the compiler and the linker!

Existing Files

Here's another common problem. You set up a `makefile` that looks like:

```

      .
      .
      lines of makefile
      .
      .
print:
      lpr $(SRCS)
      .
      .
      lines of makefile
      .
      .

```

You type

```
tutorial% make print
```

and you get the response:

```
'print' is up to date
```

instead of printing anything. The solution: there is a file called `print` in your current directory.

3.7. Making Archive Libraries

`make` provides a mechanism for referring to members of archive `ar`-style libraries. You can name a member of an object library as:

```
library-name (object-name .o)
```

or

```
library-name ( _entry-point-name )
```

The first form refers to an object name within a library. The second form refers to an *entry point* of an object file within a library. `make` searches the library to locate the entry point and then translates it to the correct object file name.

`make` has a rule for building libraries. The handle for the rule is a `.a` suffix. Then the `.c.a` is the rule for compiling a C language source file, adding it to the library, and removing the `.o` file afterwards. The internal rules that `make` employs for the `.c.a` case are:

```

.c.a:
      $(CC) -c $(CFLAGS) $<          compile the .c file
      ar rv $@ $*.o                  add it to the library
      rm -f $*.o                      get rid of the .o file

```

3.8. Suffixes and Transformation Rules

`make` itself does not know what filename suffixes are interesting or how to transform a file with one suffix into a file with another suffix. This information is stored in an internal table that has the form of a description file. If the `-r` option is used, this table is not used.

The list of suffixes is actually the dependency list for the name `.SUFFIXES`; `make` looks for a file with any of the suffixes on the list. If such a file exists, and if there is a transformation rule for that combination, `make` acts as described earlier. The transformation rule names are the concatenation of the two suffixes. The name of the rule to transform a `.c` file to a `.o` file is thus `.c.o`. If the rule is present and no explicit command sequence has been given in the user's description files, the command sequence for the rule `.c.o` is used. If a command is generated by using one of these suffixing rules, the macro `$*` is given the value of the stem (everything but the suffix) of the name of the file to be made, and the macro `$<` is the name of the dependent that caused the action.

Null Suffix

If you have many programs that are made from a single source file it is tedious to maintain an object of such files. `make` supports single suffix rules (null suffix). Suppose you have a single program called `buzz` that you maintain from a single source file `buzz.c`. You can maintain `buzz` by a `makefile` entry that looks like this:

```
.c:
    $(CC) $(CFLAGS) $(LDFLAGS) $< -o $@
```

In fact, `make` defines the `.c` rule internally so that no `makefile` is even necessary. All you have to do is type

```
tutorial% make buzz
```

and `make` will do the correct thing.

Standard Suffix List

The figures below show the rules used by `make`'s standard built-in suffix list.

NOTE *The order of the suffix list is significant, since it is scanned from left to right, and the first name that is formed that has both a file and a rule associated with it is used. If new names are to be appended, the user can just add an entry for `.SUFFIXES` in his own description file; the dependents are added to the usual list. A `.SUFFIXES` line without any dependents deletes the current list. It is necessary to clear the current list if the order of names is to be changed.*

Figure 3-1 *Single Suffix Rules*

```
.c:
    tab $(CC) $(CFLAGS) $(LDFLAGS) $< -o $@
.c~:
    tab $(GET) -G$*.c $(GFLAGS) $<
    tab $(CC) $(CFLAGS) $(LDFLAGS) $*.c -o $*
.p:
    tab $(PC) $(PFLAGS) $(LDFLAGS) $< -o $@
.p~:
```

```

        tab $(GET) -G$*.p $(GFLAGS) $<
        tab $(PC) $(PFLAGS) $(LDFLAGS) $*.p -o $*
.f:
        tab $(FC) $(FFLAGS) $(LDFLAGS) $< -o $@
.f~:
        tab $(GET) -G$*.f $(GFLAGS) $<
        tab $(FC) $(FFLAGS) $(LDFLAGS) $*.f -o $*
.F:
        tab $(FC) $(FFLAGS) $(LDFLAGS) $< -o $@
.F~:
        tab $(GET) -G$*.F $(GFLAGS) $<
        tab $(FC) $(FFLAGS) $(LDFLAGS) $*.F -o $*
.r:
        tab $(FC) $(FFLAGS) $(LDFLAGS) $< -o $@
.r~:
        tab $(GET) -G$*.r $(GFLAGS) $<
        tab $(FC) $(FFLAGS) $(LDFLAGS) $*.r -o $*
.sh:
        tab cp $< $@; chmod +x $@
.sh~:
        tab $(GET) -G$*.sh $(GFLAGS) $<
        tab cp $*.sh $*; chmod +x $@

```

Figure 3-2 Double Suffix Rules

```

.c.o:
        tab $(CC) $(CFLAGS) -c $<
.c~.o:
        tab $(GET) -G$*.c $(GFLAGS) $<
        tab $(CC) $(CFLAGS) -c $*.c
.c~.c:
        tab $(GET) -G$*.c $(GFLAGS) $<
.p.o:
        tab $(PC) $(PFLAGS) -c $<
.p~.o:
        tab $(GET) -G$*.p $(GFLAGS) $<
        tab $(PC) $(PFLAGS) -c $*.p
.p~.p:
        tab $(GET) -G$*.p $(GFLAGS) $<
.f.o:
        tab $(FC) $(FFLAGS) -c $<
.f~.o:
        tab $(GET) -G$*.f $(GFLAGS) $<
        tab $(FC) $(FFLAGS) -c $*.f
.f~.f:
        tab $(GET) -G$*.f $(GFLAGS) $<
.F.o:
        tab $(FC) $(FFLAGS) -c $<
.F~.o:
        tab $(GET) -G$*.F $(GFLAGS) $<
        tab $(FC) $(FFLAGS) -c $*.F
.F~.F:
        tab $(GET) -G$*.F $(GFLAGS) $<

```

```

.r.o:
    tab $(FC) $(FFLAGS) -c $<
.r~.o:
    tab $(GET) -G$*.r $(GFLAGS) $<
    tab $(FC) $(FFLAGS) -c $*.r
.r~.r:
    tab $(GET) -G$*.r $(GFLAGS) $<
.s.o:
    tab $(AS) $(ASFLAGS) -o $@ $<
.s~.o:
    tab $(GET) -G$*.s $(GFLAGS) $<
    tab $(AS) $(ASFLAGS) -o $*.o $*.s
.s~.s:
    tab $(GET) -G$*.s $(GFLAGS) $<
.y.o:
    tab $(YACC) $(YFLAGS) $<
    tab $(CC) $(CFLAGS) -c y.tab.c
    tab rm y.tab.c
    tab mv y.tab.o $@
.y~.o:
    tab $(GET) -G$*.y $(GFLAGS) $<
    tab $(YACC) $(YFLAGS) $*.y
    tab $(CC) $(CFLAGS) -c y.tab.c
    tab rm -f y.tab.c
    tab mv y.tab.o $*.o
.l.o:
    tab $(LEX) $(LFLAGS) $<
    tab $(CC) $(CFLAGS) -c lex.yy.c
    tab rm lex.yy.c
    tab mv lex.yy.o $@
.l~.o:
    tab $(GET) -G$*.l $(GFLAGS) $<
    tab $(LEX) $(LFLAGS) $*.l
    tab $(CC) $(CFLAGS) -c lex.yy.c
    tab rm -f lex.yy.c
    tab mv lex.yy.o $*.o
.y.c :
    tab $(YACC) $(YFLAGS) $<
    tab mv y.tab.c $@
.y~.c :
    tab $(GET) -G$*.y $(GFLAGS) $<
    tab $(YACC) $(YFLAGS) $*.y
    tab mv y.tab.c $*.c
.l.c :
    tab $(LEX) $(LFLAGS) $<
    tab mv lex.yy.c $@
.l~.c :
    tab $(GET) -G$*.l $(GFLAGS) $<
    tab $(LEX) $(LFLAGS) $*.l
    tab mv lex.yy.c $*.c
.c.a:
    tab $(CC) -c $(CFLAGS) $<
    tab ar rv $@ $*.o

```

```
        tab rm -f $*.o
.c~.a:
    tab $(GET) -G$*.c $(GFLAGS) $<
    tab $(CC) -c $(CFLAGS) $*.c
    tab ar rv $@ $*.o
    tab rm -f $*.o
.s~.a:
    tab $(GET) -G$*.s $(GFLAGS) $<
    tab $(AS) $(ASFLAGS) -o $*.o $*.s
    tab ar rv $@ $*.o
    tab -rm -f $*.o
.h~.h:
    tab $(GET) -G$*.h $(GFLAGS) $<
markfile.o: markfile
    tab A=@;echo \"static char _sccsid[] = \
    tab \042'grep $$A'(#)' markfile'\042;\> markfile.c
    tab cc -c markfile.c
    tab rm -f markfile.c
```

Source Code Control System

Source Code Control System	71
4.1. Learning the Lingo	74
S-file	74
Deltas	74
SIDs (version numbers)	74
Id keywords	74
4.2. Creating SCCS Database Files with <code>sccs create</code>	75
4.3. Retrieving Files for Compilation with <code>sccs get</code>	75
4.4. Changing Files (Creating Deltas)	76
Retrieving a File for Editing with <code>sccs edit</code>	76
Merging Changes Back Into the s-file with <code>sccs delta</code>	76
When to Make Deltas	77
Finding Out What's Going On with <code>sccs info</code>	77
ID keywords	77
Finding Out What Versions Are Being Used with <code>sccs</code> what	78
Where to Put Id Keywords	78
Keeping SIDs Consistent Across Files	78
Creating New Releases	79
4.5. Restoring Old Versions	79
Reverting to Old Versions	79
Selectively Deleting Old Deltas	79
4.6. Auditing Changes	80
Displaying Delta Comments with <code>sccs prt</code>	80

Finding Why Lines Were Inserted	80
Discovering What Changes You Have Made with <code>sccs</code> <code>diffs</code>	81
4.7. Shorthand Notations	81
Making a Delta and Getting a File with <code>sccs delget</code>	81
Replacing a Delta with the <code>sccs fix</code>	81
Backing Out of an Edit with <code>sccs unedit</code>	82
Working From Other Directories with the <code>d</code> Flag	82
4.8. Using SCCS on a Project	82
4.9. Saving Yourself	82
Recovering a Munged Edit File	82
Restoring the <code>s</code> -file	83
4.10. Managing SCCS Files with <code>sccs admin</code>	83
4.11. Maintaining Different Versions (Branches)	83
Creating a Branch	84
Getting From a Branch	84
Merging a Branch Back into the Main Trunk	84
A More Detailed Example	84
A Warning	85
4.12. Using SCCS with <code>make</code>	85
Maintaining Single Programs	86
Maintaining A Library	86
Maintaining A Large Program	87
4.13. SCCS Quick Reference	89
Commands	89
Id Keywords	90

Source Code Control System

The Source Code Control System (SCCS) is a tool for controlling changes to text files (typically, the source code and documentation of software systems).

You can think of SCCS as a custodian of files. With SCCS you can:

- Store, update, and retrieve any version of a text file.
- Control updating privileges to that file.
- Identify the version of a retrieved file.
- Record who made each change, when and where it was made, and why.

These custodial and recording functions are important in environments where programs and documentation undergo frequent changes (due to maintenance and/or enhancement work), because regenerating an unrevised version of a program or document is often desirable. Obviously, this could be done by keeping copies (on paper or other media), but this quickly becomes unmanageable and wasteful as the number of programs and documents increases. SCCS provides an attractive alternative to stockpiling multiple versions of the same text, because it stores only the original file and subsequent sets of *changes* on disk.

High-Level and Low-Level SCCS

There are two major divisions of SCCS:

- The `sccs` command itself is a high-level ‘user-friendly’ front end that provides an interface to a collection of tools for manipulating SCCS files. Basically you can type

```
tutorial% sccs do something
```

where *do something* is the operation you want to perform. In general, users can get by using the facilities provided by the `sccs` command, as described in this chapter. The individual SCCS tools are incredibly hard to use, but they do provide extremely close control over the SCCS database files.

- The individual SCCS commands are a collection of programs for manipulating the SCCS database files. Although the `sccs` front end command normally abstracts the most common operations you might want to do, there may be times when it is necessary to use the raw facilities of the SCCS commands themselves — these commands are described in appendix A, which gives a deeper description of how to use SCCS. Of particular interest are the

numbering of branches, the *l-file*, which gives a description of what deltas were used on a `sccs get`, and certain other SCCS commands.

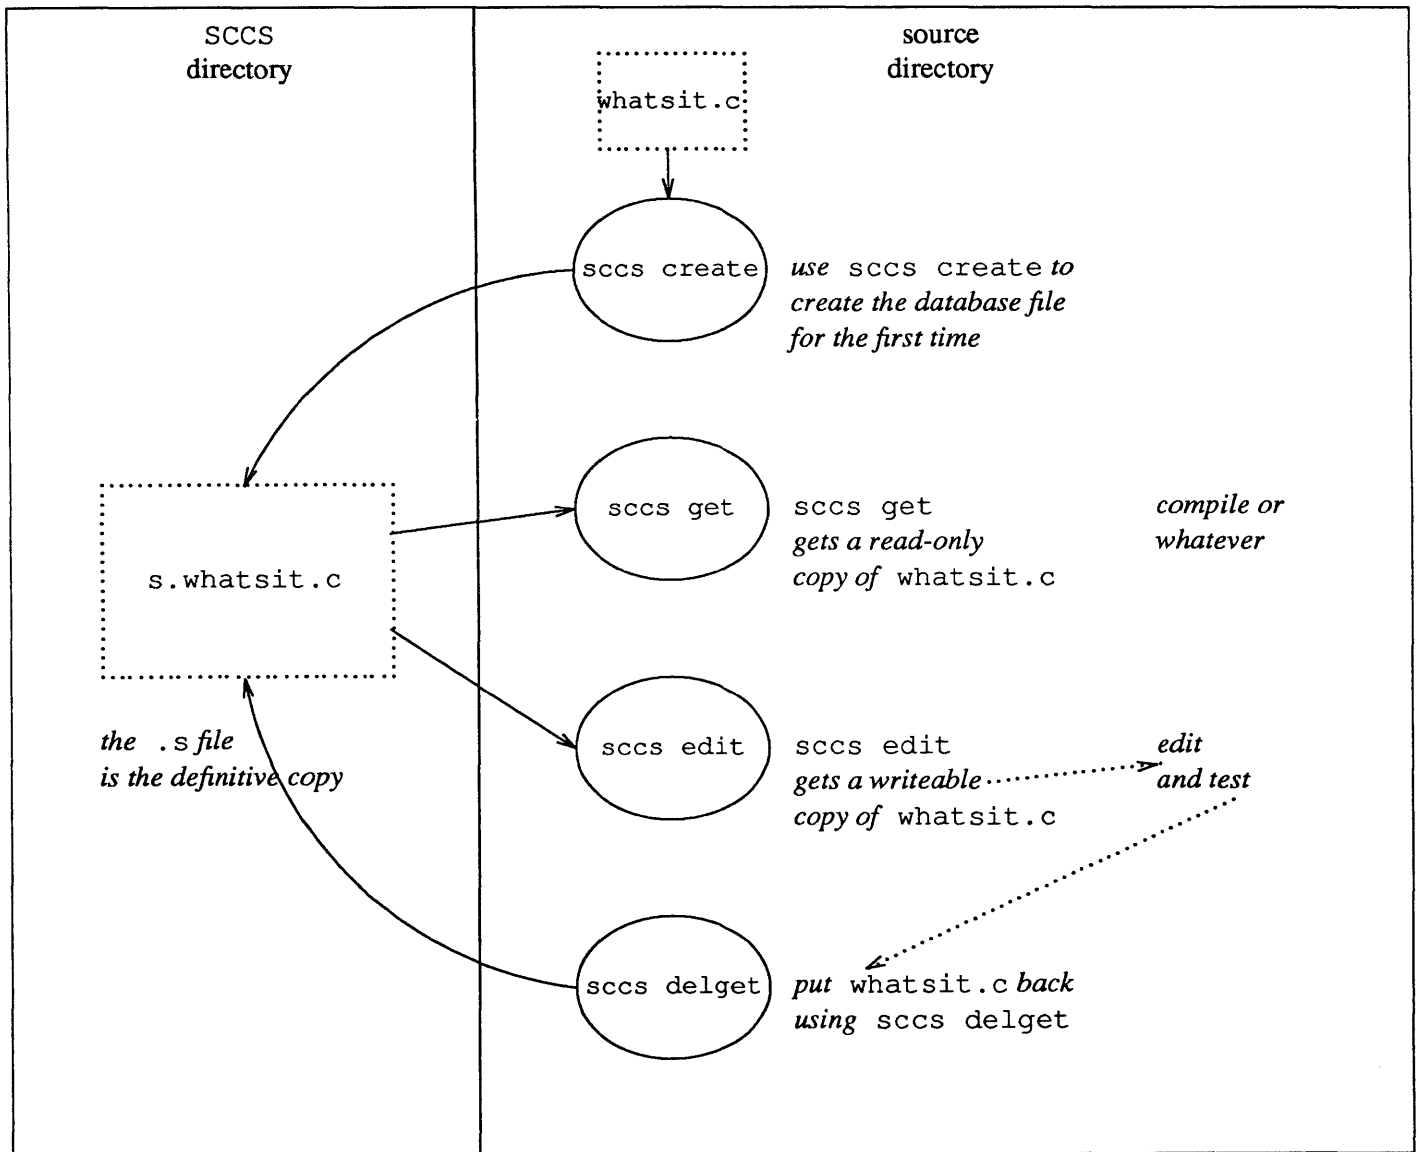
Conventions

Throughout this chapter, we assume that you are using the C-Shell on a system called 'tutorial', and so the hostname is shown followed by the `%` sign prompt in the examples. What you type is shown in **bold typewriter text like this**, and the system's responses are shown in ordinary typewriter text, like this:

```
tutorial% sccs get prog.c
1.1
87 lines
tutorial%
```

All versions of your source file, plus the log and other information, are kept in a file called the *s-file*. The illustration below shows the four basic operations that you do with SCCS.

Figure 4-1 Basic SCCS Operations



As the picture illustrates, there are four major operations that can be performed on the *s-file* :

- *create* the *s-file* in the very first place.
- *get* a read-only copy of the *s-file* . This operation retrieves a version of the file from the *s-file* . By default, the latest version is retrieved. This read-only copy is intended for compilation, printing, or whatever — it is specifically NOT intended to be edited or changed in any way — any changes made to a file retrieved in this way will probably be lost.
- *Get* a file for editing. This operation also retrieves a version of the file from the *s-file* , but this file is intended to be edited and then incorporated back into the *s-file* . Only one person may be editing a file at one time.

- Merge any changes made back into the *s-file*. This is the companion operation to the previous operation. A new version number is assigned, and comments are saved explaining why this change was made.

Understand that the *s-file* is the 'real' instance of whatever file it is you are working with. The copy you get from the SCCS database by using a `sccs get` or a `sccs edit` command is a *copy*, and should be considered ephemeral.

4.1. Learning the Lingo

There are a number of terms that are worth learning before we go any farther.

S-file

The *s-file* is a single file that holds all the different versions of your file. The *s-file* contains only the the original version and differences between versions, rather than the entire text of the new version. This saves disk space and allows selective changes to be removed later. Also included in the *s-file* is some header information for each version, including the comments given by the person who created the version explaining why the changes were made.

Deltas

Each set of changes to the *s-file* — which is approximately, but not exactly, equivalent to a version of the file — is called a *delta*. Although technically a delta only includes the *changes* made, in practice it is usual for each delta to be made with respect to all the deltas that have occurred before¹. However, it is possible to get a version of the file that has selected deltas removed out of the middle of the list of changes — equivalent to removing your later changes.

SIDs (version numbers)

An SID — SCCS-Id — is a number that represents a delta. This is normally a two-part number consisting of a 'release' number and a 'level' number. Normally the release number stays the same. However, it is possible to move into a new release if some major change is being made.

Since all past deltas are normally applied, the SID of the final delta applied can be used to represent a version number of the file as a whole.

Id keywords

When you get a version of a file with intent to compile and install it — that is, something other than edit it — some special keywords that are part of the text of the file are expanded in-line by SCCS. These *Id Keywords* can be used to include the current version number or other information into the file. All id keywords are of the form `%x%`, where *x* is an upper case letter. For example, `%I%` produces the SID of the latest delta applied, `%W%` includes the module name, SID, and a mark that makes it findable by a program, and `%G%` results in the date the latest delta was applied. There are many others, most of which are of dubious value.

When you get a file for editing, the id keywords are not expanded; this is so that after you put them back in to the *s-file*, they will be expanded automatically on each new version. But notice: if you were to get them expanded accidentally, your file would appear to be the same version forever more, which would of course defeat the purpose. Also, if you should install a version of the program without expanding the id keywords, it will be impossible to tell what version it is (since

¹ This matches normal usage, where the previous changes are not saved at all, so all changes are automatically based on all other changes that have happened through history.

4.2. Creating SCCS Database Files with `sccs create`

all it will have is `%W%` or whatever).

To put a bunch of source files into SCCS format, you do the following things:

- Make the SCCS subdirectory if it isn't there already:

```
tutorial% mkdir SCCS    Note that SCCS is upper-case
tutorial%
```

- Then use the `sccs create` command to actually create the SCCS database files for all the source files you have. Suppose that you want to have all your `.c` and `.h` files under SCCS control:

```
tutorial% sccs create *. [ch]
           lots of messages from SCCS here
tutorial%
```

For each *file* you have, the `sccs create` command does the following things for you:

- creates* a file called *s.file* in the SCCS subdirectory,
- renames* each *file* by placing a comma in front of the name, so that you end up with files of the form *,file*.
- gets* a read-only copy of each *file* by using the `sccs get` command, described later on.

When you are convinced that SCCS has correctly created the *s-file*s, you should remove the files whose names start with commas.

If you want to have id keywords in the files, it is best to put them in before you create the *s-files*. If you do not, `create` will print

```
No Id Keywords (cm7)
```

which is a warning message only.

4.3. Retrieving Files for Compilation with `sccs get`

To get a copy of the latest version of a file, run

```
tutorial% sccs get prog.c
```

SCCS will respond:

```
1.1
87 lines
```

meaning that version 1.1 has been retrieved² and that it has 87 lines. The file `prog.c` is created in the current directory — it is created read-only to remind

² Actually, the `sm` of the final delta applied was 1.1.

you that you are not supposed to change it.

This copy of the file should not be changed, since SCCS is unable to merge the changes back into the *s-file*. If you do make changes, they will be lost the next time someone does a `sccs get`.

4.4. Changing Files (Creating Deltas)

To change a version of a file, you must obtain a copy of the file that can be edited. You obtain such a copy using `sccs get` as shown below. Having made the changes and satisfied yourself that the changes are correct, you can then merge the changes back into the SCCS database file using `sccs delta` also shown below.

Retrieving a File for Editing with `sccs edit`

To edit a source file, you must first get it, requesting permission to edit it³. The response will be the same as with `sccs get` except that it also says that a new delta is being created:

```
tutorial% sccs edit prog.c
New delta 1.2
```

You then edit it, using a text editor:

```
tutorial% vi prog.c
```

database file using

Merging Changes Back Into the *s-file* with `sccs delta`

When the desired changes have been made, you can put your changes into the SCCS file using the *delta* command:

```
tutorial% sccs delta prog.c
```

Delta prompts you for 'comments?' before merging the changes in. At this prompt you should type a one-line description of what the changes mean (more lines can be entered by ending each line except the last with a backslash). *Delta* then types:

```
1.2
5 inserted
3 deleted
84 unchanged
```

saying that delta 1.2 was created, and it inserted five lines, removed three lines, and left 84 lines unchanged⁴. The `prog.c` file is then removed; it can be retrieved using `sccs get`.

³ The `sccs edit` command is equivalent to using the `-e` option to `sccs get`, as:

```
tutorial% sccs get -e prog.c
```

Keep this in mind when reading other documentation.

⁴ Changes to a line are counted as a line deleted and a line inserted.

When to Make Deltas

It is probably unwise to make a delta before every recompilation or test; otherwise, you tend to get a lot of deltas with comments like 'fixed compilation problem in previous delta' or 'fixed botch in 1.3'. However, it is very important to delta everything before installing a module for general use. A good technique is to edit the files you need, make all necessary changes and tests, compiling and editing as often as necessary without making deltas. When you are satisfied that you have a working version, delta everything being edited, re-get them, and recompile everything.

Finding Out What's Going On with `sccs info`

To find out what files are being edited, type:

```
tutorial% sccs info
```

to display a list of all the files being edited and other information — such as the name of the user who did the edit. Also, the command:

```
tutorial% sccs check
```

is nearly equivalent to the `info` command, except that it is silent if nothing is being edited, and returns non zero exit status if anything is being edited. It can thus be used in an 'install' entry in a makefile to abort the install if anything has not been properly delta'ed.

If you know that everything being edited should be delta'ed, you can use:

```
tutorial% sccs delta `sccs tell`
```

The `tell` command is similar to `info` except that only the names of files being edited are output, one per line.

All of these commands take a `-b` option to ignore 'branches' (alternate versions, described later) and the `-u` option to give only files being edited by you. The `-u` option takes an optional `user` argument, giving only files being edited by that user. For example:

```
tutorial% sccs info -ujohn
```

gives a listing of files being edited by john.

ID keywords

Id keywords can be inserted into your file that will be expanded automatically by `sccs get`. For example, a line such as:

```
static char SccsId[ ] = "%W%\t%G%";
```

will be replaced with something like:

```
static char SccsId[ ] = "@(#)prog.c 1.2 08/29/80";
```

This tells you the name and version of the source file and the time the delta was created. The string '@(#)' is a special string which signals the beginning of an

Finding Out What Versions Are Being Used with `sccs what`

SCCS Id keyword.

To find out what version of a program is being run, use:

```
tutorial% sccs what prog.c /usr/bin/prog
```

which will print all strings it finds that begin with '@ (#)'. This works on all types of files, including binaries and libraries. For example, the above command will output something like:

```
prog.c:
  prog.c 1.2 08/29/80
/usr/bin/prog:
  prog.c 1.1 02/05/79
```

From this one can see that the source in *prog.c* will not compile into the same version as the binary in */usr/bin/prog*.

Where to Put Id Keywords

ID keywords can be inserted anywhere, including in comments, but Id keywords that are compiled into the object module are especially useful, since they let you find out what version of the object is being run. However, there is a cost: data space is used up to store the keywords.

When you put id keywords into header files, it is important that you assign them to different variables. For example, you might use:

```
static char AccessSid[ ] = "%W% %G%";
```

in the file *access.h* and:

```
static char OpsysSid[ ] = "%W% %G%";
```

in the file *opsys.h*. Otherwise, you will get compilation errors because 'SccsId' is redefined. The problem with this is that if the header file is included by many modules that are loaded together, the version number of that header file is included in the object module many times; you may find it more to your taste to put id keywords in header files in comments.

Keeping SIDs Consistent Across Files

With some care, it is possible to keep the SID's consistent in multi-file systems. The trick here is to always `sccs edit` all files at once. The changes can then be made to whatever files are necessary and then all files (even those not changed) are redelta'ed. This can be done fairly easily by just specifying the name of the directory that the SCCS files are in:

```
tutorial% sccs edit SCCS
```

which will `sccs edit` all files in that directory. To make the delta, use:

```
tutorial% sccs delta SCCS
```

You will be prompted for comments only once.

Creating New Releases

When you want to create a new release of a program, you can specify the release number you want to create on the `sccs edit` command. For example:

```
tutorial% sccs edit -r2 prog.c
```

will put the next delta in release two (that is, it will be numbered 2.1). Future deltas will automatically be in release two. To change the release number of an entire system, use:

```
tutorial% sccs edit -r2 SCCS
```

4.5. Restoring Old Versions Reverting to Old Versions

Suppose that after delta 1.2 was stable you made and released a delta 1.3. But this introduced a bug, so you made a delta 1.4 to correct it. But 1.4 was still buggy, and you decided you wanted to go back to the old version. You could revert to delta 1.2 by choosing the SID in a get:

```
tutorial% sccs get -r1.2 prog.c
```

This will produce a version of `prog.c` that is delta 1.2 that can be reinstalled so that work can proceed.

In some cases you don't know what the SID of the delta you want is. However, you can revert to the version of the program that was running as of a certain date by using the `-c` (cutoff) option. For example,

```
tutorial% sccs get -c800722120000 prog.c
```

retrieves whatever version was current as of July 22, 1980 at 12:00 noon. Trailing components can be stripped off (defaulting to their highest legal value), and punctuation can be inserted in the obvious places; for example, the above line could be equivalently stated as:

```
tutorial% sccs get -c"80/07/22 12:00:00" prog.c
```

Selectively Deleting Old Deltas

Suppose that you later decided that you liked the changes in delta 1.4, but that delta 1.3 should be removed. You could do this by *excluding* delta 1.3:

```
tutorial% sccs edit -x1.3 prog.c
```

When delta 1.5 is made, it will include the changes made in delta 1.4, but will exclude the changes made in delta 1.3. You can exclude a range of deltas using a dash. For example, if you want to get rid of 1.3 and 1.4 you can use:

```
tutorial% sccs edit -x1.3-1.4 prog.c
```

which will exclude all deltas from 1.3 through 1.4. Alternatively,

```
tutorial% sccs edit -x1.3-1 prog.c
```

will exclude a range of deltas from 1.3 to the current highest delta in release 1.

In certain cases when using `-x` (or `-i` — see below) there will be conflicts between versions; for example, it may be necessary to both include and delete a particular line. If this happens, SCCS always displays a message telling the range of lines affected; these lines should then be examined very carefully to see if the version SCCS got is ok.

Since each delta (in the sense of 'a set of changes') can be excluded at will, it is most useful to put each semantically distinct change into its own delta.

4.6. Auditing Changes

Displaying Delta Comments with `sccs prt`

When you created a delta, you presumably gave a reason for the delta to the 'comments?' prompt. To display these comments later, use:

```
tutorial% sccs prt prog.c
```

which produces a report for each delta of the SID, time and date of creation, user who created the delta, number of lines inserted, deleted, and unchanged, and the comments associated with the delta. For example, the output of the above command might be:

```
tutorial% sccs prt prog.c
D 1.2  80/08/29 12:35:31  bill    2   100005/00003/00084
removed "-q" option

D 1.1  79/02/05 00:19:31  eric    1   000087/00000/00000
date and time created 80/06/10 00:19:31 by eric
```

Finding Why Lines Were Inserted

To find out why you inserted lines, you can get a copy of the file with each line preceded by the SID that created it:

```
tutorial% sccs get -m prog.c
```

You can then find out what changes were made by this delta by printing the comments using `prt`.

To find out what lines are associated with a particular delta, 1.3 for instance, use:

```
tutorial% sccs get -m -p prog.c | grep '^1.3'
```

Discovering What Changes You Have Made with `sccs diffs`

The `-p` option makes SCCS output the generated source to the standard output rather than to a file.

When you are editing a file, you can find out what changes you have made using:

```
tutorial% sccs diffs prog.c
```

Most of the “diff” options can be used. To pass the `-c` option, use `-C`.

To compare two versions that are in deltas, use:

```
tutorial% sccs sccsdiff -r1.3 -r1.6 prog.c
```

to see the differences between delta 1.3 and delta 1.6.

4.7. Shorthand Notations

Making a Delta and Getting a File with `sccs delget`

There are several sequences of commands that are used frequently. *Sccs* tries to make it easy to do these.

A frequent requirement is to make a delta of some file and then get that file. This is done by using

```
tutorial% sccs delget prog.c
```

which is entirely equivalent to:

```
tutorial% sccs delta prog.c
tutorial% sccs get prog.c
```

except that if an error occurs while making a delta of *any* of the files, none of them will be gotten. The `sccs deleedit` command is equivalent to `sccs delget` except that the `sccs edit` command is used instead of the `sccs get` command.

Replacing a Delta with the `sccs fix`

Frequently, there are small bugs in deltas, for instance, compilation errors, for which there is no reason to maintain an audit trail. To *replace* a delta, use:

```
tutorial% sccs fix -r1.4 prog.c
```

This gets a copy of delta 1.4 of `prog.c` for you to edit and then deletes delta 1.4 from the SCCS file. When you do a delta of `prog.c`, it will be delta 1.4 again. The `-r` option must be specified, and the delta that is specified must be a leaf delta, that is, no other deltas may have been made subsequent to the creation of that delta.

Backing Out of an Edit with `sccs unedit`

If you found you edited a file that you did not want to edit, you can back out by using:

```
tutorial% sccs unedit prog.c
```

Working From Other Directories with the `-d` Flag

If you are working on a project where the SCCS code is in a directory somewhere else, you may be able to simplify things by using a shell alias. For example, the alias:

```
alias syssccs sccs -d/usr/src
```

will allow you to issue commands such as:

```
syssccs edit cmd/who.c
```

which will look for the file `'usr/src/cmd/SCCS/who.c'`. The file `'who.c'` is always created in your current directory regardless of the value of the `-d` option.

4.8. Using SCCS on a Project

Working on a project with several people has its own set of special problems. The main problem occurs when two people modify a file at the same time. SCCS prevents this by locking an *s-file* while it is being edited.

As a result, files should not be reserved for editing unless they are actually being edited at the time, since this will prevent other people on the project from making necessary changes. For example, a good scenario for working might be:

```
tutorial% sccs edit a.c g.c t.c  
tutorial% vi a.c g.c t.c  
# do testing of the (experimental) version  
tutorial% sccs delget a.c g.c t.c  
tutorial% sccs info  
# should respond "Nothing being edited"  
tutorial% make install
```

As a general rule, all source files should be delta'ed before installing the program for general use. This will ensure that it is possible to restore any version in use at any time.

4.9. Saving Yourself Recovering a Munged Edit File

Sometimes you may find that you have destroyed or trashed a file that you were trying to edit⁵. Unfortunately, you can't just remove it and re-`sccs edit` it; SCCS keeps track of the fact that someone is trying to edit it, so it won't let you do it again. Neither can you just get it using `sccs get`, since that would expand the Id keywords. Instead, you can say:

```
tutorial% sccs get -k prog.c
```

⁵ Or given up and decided to start over.

This will not expand the Id keywords, so it is safe to do a delta with it.

Alternatively, you can *unedit* and `sccs edit` the file.

Restoring the s-file

In particularly bad circumstances, the SCCS file itself may get munged. The most common way this happens is that it gets edited. Since SCCS keeps a checksum, you will get errors every time you read the file. To fix this checksum, use:

```
tutorial% sccs admin -z prog.c
```

4.10. Managing SCCS Files with `sccs admin`

There are a number of parameters that can be set using the *admin* command. The most interesting of these are flags. Flags can be added by using the `-f` option. For example:

```
tutorial% sccs admin -fd1 prog.c
```

sets the 'd' flag to the value '1'. This flag can be deleted by using:

```
tutorial% sccs admin -dd prog.c
```

The most useful flags are:

b Allow branches to be made using the `-b` option to `sccs edit`.

dSID

Default SID to be used on a `sccs get` or `sccs edit`. If this is just a release number it constrains the version to a particular release only.

i Give a fatal error if there are no Id keywords in a file. This is useful to guarantee that a version of the file does not get merged into the *s-file* that has the Id keywords inserted as constants instead of internal forms.

y The 'type' of the module. Actually, the value of this flag is unused by SCCS except that it replaces the `%Y%` keyword.

-tfile

store descriptive text from *file* in the SCCS file. This descriptive text might be the documentation or a design and implementation document. Using the `-t` option ensures that if the SCCS file is passed on to someone else, the documentation will go along with it. If *file* is omitted, the descriptive text is deleted. To see the descriptive text, use `prt -t`.

The *admin* command can be used safely any number of times on files. A file need not be gotten for *admin* to work.

4.11. Maintaining Different Versions (Branches)

Sometimes it is convenient to maintain an experimental version of a program for an extended period while normal maintenance continues on the version in production. This can be done using a 'branch'. Normally deltas continue in a straight line, each depending on the delta before. Creating a branch 'forks off' a version of the program.

The ability to create branches must be enabled in advance using:

```
tutorial% sccs admin -fb prog.c
```

The `-fb` option can be specified when the SCCS file is first created.

Creating a Branch

To create a branch, use:

```
tutorial% sccs edit -b prog.c
```

This will create a branch with (for example) SID 1.5.1.1. The deltas for this version will be numbered 1.5.1.*n*.

Getting From a Branch

Deltas in a branch are normally not included when you do a `get`. To get these versions, you will have to say:

```
tutorial% sccs get -r1.5.1 prog.c
```

Merging a Branch Back into the Main Trunk

At some point you will have finished the experiment, and if it was successful you will want to incorporate it into the released version. But in the meantime someone may have created a delta 1.6 that you don't want to lose. The commands:

```
tutorial% sccs edit -i1.5.1.1-1.5.1 prog.c  
tutorial% sccs delta prog.c
```

will merge all of your changes into the release system. If some of the changes conflict, `get` will print an error. The generated result should be carefully examined before the delta is made.

A More Detailed Example

The following technique might be used to maintain a different version of a program. First, create a directory to contain the new version:

```
tutorial% mkdir ../newxyz  
tutorial% cd ../newxyz
```

Edit a copy of the program on a branch:

```
tutorial% sccs -d../xyz edit -b prog.c
```

When using the old version, be sure to use the `-b` option to *info*, *check*, *tell*, and *clean* to avoid confusion. For example, use:

```
tutorial% sccs info -b
```

when in the 'xyz' directory.

If you want to save a copy of the program (still on the branch) back in the *s-file*, you can use:

```
tutorial% sccs -d../xyz deledit prog.c
```

which will do a delta on the branch and reedit it for you.

When the experiment is complete, merge it back into the *s-file* using delta:

```
tutorial% sccs -d../xyz delta prog.c
```

At this point you must decide whether this version should be merged back into the trunk, that is, the default version, which may have undergone changes. If so, it can be merged using the `-i` option to `sccs edit` as described above.

A Warning

Branches should be kept to a minimum. After the first branch from the trunk, SID's are assigned rather haphazardly, and the structure gets complex fast.

4.12. Using sccs with make

SCCS and make can be made to work together with a little care. A few sample makefiles for common applications are shown below.

There are a few basic entries that every *Makefile* ought to have. These are:

- a.out (or whatever the *Makefile* generates). This entry regenerates a program. If the *Makefile* regenerates many things, this should be called 'all' and should in turn have dependencies on everything the *Makefile* can generate.
- install Moves the objects to their final resting place, doing any special *chmod*'s or *ranlib*'s as appropriate.
- sources Creates all the source files from SCCS files.
- clean Removes all unwanted files from the directory.
- print Prints the contents of the directory.

The examples shown below are only partial examples, and may omit some of these entries when they are deemed to be obvious.

The *clean* entry should not remove files that can be regenerated from the SCCS files. It is sufficiently important to have the source files around at all times that the only time they should be removed is when the directory is being mothballed. To do this, the command:

```
tutorial% sccs clean
```

can be used. This removes all files for which an *s-file* exists, but which are not being edited.

Maintaining Single Programs

Frequently there are directories with several largely unrelated programs (such as simple commands). These can be put into a single *Makefile* :

```
LDLFLAGS= -i -s
prog: prog.o
    $(CC) $(LDLFLAGS) -o prog prog.o
prog.o: prog.c prog.h
example: example.o
    $(CC) $(LDLFLAGS) -o example example.o
example.o: example.c
.DEFAULT:
    sccs get $<
```

The trick here is that the `.DEFAULT` rule is called every time something is needed that does not exist, and no other rule exists to make it. The explicit dependency of the `.o` file on the `.c` file is important. Another way of doing the same thing is:

```
SRCS= prog.c prog.h example.c
LDLFLAGS= -i -s
prog: prog.o
    $(CC) $(LDLFLAGS) -o prog prog.o
prog.o: prog.h
example: example.o
    $(CC) $(LDLFLAGS) -o example example.o
sources: $(SRCS)
$(SRCS):
    sccs get $@
```

There are a couple of advantages to this approach: (1) the explicit dependencies of the `.o` on the `.c` files are not needed, (2) there is an entry called "sources" so if you want to get all the sources you can just say 'make sources' and (3) the makefile is less likely to do confusing things since it won't try to `sccs get` things that do not exist.

Maintaining A Library

Libraries that are largely static are best updated using explicit commands, since *make* doesn't know about updating them properly. However, libraries that are in the process of being developed can be handled quite adequately. The problem is that the `.o` files have to be kept separate from the library, as well as in the library.

```

# configuration information
OBJS=  a.o b.o c.o d.o
SRCS=  a.c b.c c.c d.s x.h y.h z.h
TARG=  /usr/lib
# programs
GET=   sccs get
REL=
AR=   -ar
RANLIB= ranlib
lib.a: $(OBJS)
        $(AR) rvu lib.a $(OBJS)
        $(RANLIB) lib.a
install: lib.a
        sccs check
        cp lib.a $(TARG)/lib.a
        $(RANLIB) $(TARG)/lib.a
sources: $(SRCS)
$(SRCS):
        $(GET) $(REL) $@
print: sources
        pr *.h *.c[cs]
clean:
        rm -f *.o
        rm -f core a.out $(LIB)

```

The '\$(REL)' in the get can be used to get old versions easily; for example:

```
tutorial% make b.o REL=-r1.3
```

The *install* entry includes the line `sccs check` before anything else. This guarantees that all the *s-file*'s are up-to-date (that is, nothing is being edited), and will abort the *make* if this condition is not met.

Maintaining A Large Program

```

OBJS=  a.o b.o c.o d.o
SRCS=  a.c b.c y.c d.s x.h y.h z.h
GET=   sccs get
REL=
a.out: $(OBJS)
        $(CC) $(LDFLAGS) $(OBJS) $(LIBS)
sources: $(SRCS)
$(SRCS):
        $(GET) $(REL) $@

```

The *print* and *clean* entries are identical to the previous case. This *Makefile* requires copies of the source and object files to be kept during development. It is probably also wise to include lines of the form:

```
a.o: x.h y.h
b.o: z.h
c.o: x.h y.h z.h
z.h: x.h
```

so that modules will be recompiled if header files change.

Since *make* does not do transitive closure on dependencies, you may find in some *Makefile* s lines like:

```
z.h: x.h
    touch z.h
```

This would be used in cases where file *z.h* has a line:

```
#include "x.h"
```

The *touch* command brings the modification date of *z.h* in line with the modification date of *x.h*. When you have a *Makefile* such as the above, the *touch* command can be removed completely; the equivalent effect will be achieved by doing an automatic `sccs get` on *z.h*.

4.13. SCCS Quick Reference

Commands	This list is not exhaustive; for more options see appendix A of this manual.
<code>sccs get</code>	Gets files for compilation (not for editing). Id keywords are expanded. <ul style="list-style-type: none"> <code>-rSID</code> Version to get. <code>-p</code> Send to standard output rather than to the actual file. <code>-k</code> Don't expand id keywords. <code>-i list</code> List of deltas to include. <code>-x list</code> List of deltas to exclude. <code>-m</code> Precede each line with SID of creating delta. <code>-c date</code> Don't apply any deltas created after <i>date</i>.
<code>sccs edit</code>	Gets files for editing. Id keywords are not expanded. Should be matched with a <code>delta</code> command. <ul style="list-style-type: none"> <code>-rSID</code> Same as for <code>sccs get</code>. If <i>SID</i> specifies a release that does not yet exist, the highest numbered delta is retrieved and the new delta is numbered with <i>SID</i>. <code>-b</code> Create a branch. <code>-i list</code> Same as for <code>sccs get</code>. <code>-x list</code> Same as for <code>sccs get</code>.
<code>sccs delta</code>	Merge a file gotten using <code>edit</code> back into the <i>s-file</i> . Collect comments about why this delta was made.
<code>sccs unedit</code>	Remove a file that has been edited previously without merging the changes into the <i>s-file</i> .
<code>sccs prt</code>	Produce a report of changes. <ul style="list-style-type: none"> <code>-t</code> Print the descriptive text. <code>-e</code> Print (nearly) everything.
<code>sccs info</code>	Give a list of all files being edited. <ul style="list-style-type: none"> <code>-b</code> Ignore branches. <code>-u[user]</code> Ignore files not being edited by <i>user</i>.
<code>sccs check</code>	Same as <code>info</code> , except that nothing is printed if nothing is being edited and exit status is returned.

<code>sccs tell</code>	Same as <code>info</code> , except that one line is produced per file being edited containing only the file name.
<code>sccs clean</code>	Remove all files that can be regenerated from the <i>s-file</i> .
<code>sccs what</code>	Find and print id keywords.
<code>sccs admin</code>	Create or set parameters on <i>s-file</i> 's. <code>-i file</code> Create, using <i>file</i> as the initial contents. <code>-z</code> Rebuild the checksum in case the file has been trashed. <code>-f flag</code> Turn on <i>flag</i> . <code>-d flag</code> Turn off (delete) <i>flag</i> . <code>-t file</code> Replace the descriptive text in the <i>s-file</i> with the contents of <i>file</i> . If <i>file</i> is omitted, the text is deleted. Useful for storing documentation or design and implementation documents to ensure they get distributed with the <i>s-file</i> . Useful flags that can be introduced via the <code>-F</code> and <code>-d</code> options are: <code>b</code> Allow branches to be made using the <code>-b</code> option to <code>edit</code> . <code>dSID</code> Default SID to be used on a <code>get</code> or <code>edit</code> . <code>i</code> Make the 'No Id Keywords' error message a fatal error rather than a warning. <code>t</code> The module 'type'; the value of this flag replaces the <code>%Y%</code> keyword.
<code>sccs fix</code>	Remove a delta and reedit it.
<code>sccs delget</code>	Do a <code>delta</code> followed by a <code>get</code> .
<code>sccs deledit</code>	Do a <code>delta</code> followed by an <code>edit</code> .
Id Keywords	<code>%Z%</code> Expands to '@(#)' for the <code>what</code> command to find. <code>%M%</code> The current module name, for example, <code>prog.c</code> . <code>%I%</code> The highest SID applied. <code>%W%</code> A shorthand for <code>%Z%%M% <tab> %I%</code> . <code>%G%</code> The date of the delta corresponding to the <code>%I%</code> keyword. <code>%R%</code> The current release number, that is, the first component of the <code>%I%</code> keyword. <code>%Y%</code> Replaced by the value of the <code>t</code> flag (set by <code>admin</code>).

Performance Analysis

Performance Analysis	93
5.1. <code>time</code> — Display Time Used by Program	93
5.2. <code>prof</code> — Generate Profile of Program	96
5.3. <code>gprof</code> — Generate Call Graph Profile of Program	98
5.4. <code>tcov</code> — Statement-level Analysis of Program	100

Performance Analysis

Tools discussed in this chapter cover facilities for timing programs and getting performance analysis data. Some tools work only with the C programming language, while others will work on modules written in any language. Performance analysis tools provide a variety of levels of analysis from very simple timing of a command down to a statement-by-statement analysis of a program. You can select which level of granularity you like depending on the amount of detail and optimization you wish to perform. Here are the performance analysis tools available from the simplest to the most detailed:

<code>time</code>	A simple command (built in to the C Shell) to display the time that a program takes. The C Shell's built in <code>time</code> command display statistics about how a command uses the system resources as well as just the raw time consumed.
<code>prof</code>	Generates a profile for the modules in a program, showing which modules are using the time.
<code>gprof</code>	Generates not only a profile as for <code>prof</code> , but also generates a <i>call graph</i> showing what modules call which, and which modules are called by other modules. The call graph can sometimes point out areas where removing calls can speed up a program.
<code>tcov</code>	Generates a detailed statement-by-statement analysis of a C program.

5.1. `time` — Display Time Used by Program

Two distinct versions of the `time` command exist in the Sun system. Here we discuss the `time` command that is built in to the C-Shell. The other `time` command is a program (in `/bin/time`) that you get when you use the Bourne Shell.

As a first example, we show the `time` command being used to display statistics on the run-time of the `index.assist` program we've used in other examples in this manual. In all the examples shown here we direct the output from `index.assist` into `/dev/null`. Here is the simplest example of using `time`:

```
tutorial% time index.assist < index.entries > /dev/null
13.5u 0.8s 0:15 92% 3+19k 19+1io 0pf+0w
tutorial%
```

Now to explain the items in the display from the `time` command above:

The 13.5u means that this program used 13.5 seconds of *user* time — time spent in the application program itself. The 0.8s means that the program spent 0.8 seconds in the *system* — this is time spent in the operating system kernel on behalf of the program. The third field is the elapsed or wallclock time for the application. The percentage figure is the percent of the user and system time as a fraction of the elapsed time. The rest of the display is of lesser interest just now and is explained in more detail below.

Effects of Optimizer on Timing

Just for the sake of interest, let's see what effect the C optimizer has on the run time of this program — we make the program with the `-O` option and see what happens:

```
tutorial% time index.assist < index.entries > /dev/null
13.1u 1.4s 0:38 37% 3+19k 19+0io 1pf+0w
tutorial%
```

What has happened here? The optimized version takes longer to run! This demonstration tells us that simple timing is not so simple after all — in a multi-tasking system there are many other factors that can effect the simple timing. Note that the user time for the program is actually slightly less — 0.4 seconds less. But, the system time and the elapsed time are very different. These timings are affected by the load on the system. If we look at the last field in the time display, note that in the unoptimized version there were zero page faults, while in the optimized version there was one page fault. This is an indication that there was other activity in the system at the time the program was run and this other activity will adversely affect the elapsed time. There are two rules you can apply to this situation:

- Run such timing tests on a quiet system late at night. Make sure that 'late at night' is not midnight when a whole bunch of `cron` daemons start up.
- Run timing tests several times and take averages.

Controlling the display from the `time` Command

The `time` command built into the C Shell has the capability of altering the information displayed under control of an environment variable. This is not true of `/bin/time` — the command you'd have to use if you were using the Bourne Shell. Here is how to set up the `time` variable to control the time display.

You can control how the C Shell times programs by setting the `time` variable in your `.login` or `.cshrc` file.

The `time` variable can be supplied with one or two values, such as `set time=3` or `set time=(3 "%E %P%)`.

Setting the `time` variable via a `set` command of the form:

```
set time=nnn
```

means that the Shell displays a resource-usage summary for any command running for more than `nnn` CPU seconds.

Control Key Letters for the time Command

The second form controls exactly what resources are displayed. The character string can be any string of text with embedded control key-letters in it. A control key-letter is a percent sign (%) followed by a single *upper-case* letter. To print a percent sign, use two percent signs in a row. Unrecognized key-letters are simply printed. The control key-letters are:

Table 5-1 *Control Key Letters for the time Command*

<i>Letter</i>	<i>Description</i>
D	Average amount of unshared data space used in Kilobytes.
E	Elapsed (wallclock) time for the command.
F	Page faults.
I	Number of block input operations.
K	Average amount of unshared stack space used in Kilobytes.
M	Maximum real memory used during execution of the process.
O	Number of block output operations.
P	Total CPU time — U (user) plus S (system) — as a percentage of E (elapsed) time.
S	Number of seconds of CPU time consumed by the kernel on behalf of the user's process.
U	Number of seconds of CPU time devoted to the user's process.
W	Number of swaps.
X	Average amount of shared memory used in Kilobytes.

Default Timing Summary

The default resource-usage summary is a line of the form:

```
uuu.uu sss.s ee:ee pp% xxx+dddk iii+ooo io mmm pf+www
```

Table 5-2 *Default Timing Summary Chart*

<i>Field</i>	<i>Description</i>
<i>uuu.u</i>	user time (U),
<i>sss.s</i>	system time (S),
<i>ee:ee</i>	elapsed time (E),
<i>pp</i>	percentage of CPU time versus elapsed time (P),
<i>xxx</i>	average shared memory in Kilobytes (X),
<i>ddd</i>	average unshared data space in Kilobytes (D),
<i>iii</i> and <i>ooo</i>	the number of block input and output operations respectively (I and O),
<i>mmm</i>	number of page faults (F)
<i>ww</i>	number of swaps (W).

C-Shell time Command versus /bin/time

One final note on the `time` commands. As mentioned previously, there are two versions of `time`: the one built in to the C-Shell as described above, and the original Bourne Shell `time` command which can be found in `/bin/time`.

The C-Shell `time` command does not time a command which is a component of a pipeline. This is what happens:

```
tutorial% echo timing a pipeline | time cat
timing a pipeline
tutorial%
```

whereas the Bourne Shell `time` command gives completely different results:

```
tutorial% echo timing a pipeline | /bin/time cat
timing a pipeline
      0.8 real          0.0 user          0.1 sys
tutorial%
```

5.2. `prof` — Generate Profile of Program

After simple timing, a *profile* of a program displays a finer level of analysis to assist in optimizing performance. Getting a profile is the next step after simple timing — more detailed analysis is provided by the *call-graph* profile and the *code coverage* tools described later in this chapter.

Taking the `index.assist` program from before as an example, let's make the program compiled for profiling. To compile a program for profiling, you use the `-p` option to the C compiler:

```
tutorial% make CFLAGS=-p
.
.
  messages from the make command
.
.
tutorial%
```

Now we can run the `index.assist` program as before. When a program is profiled, the results appear in a file called `mon.out` at the end of the run. Every time you run the program a new `mon.out` file is created, overwriting the old version. You then use the `prof` command to interpret the results of the profile:

```
tutorial% index.assist < index.entries > /dev/null
tutorial% prof index.assist
%time  cumsecs  #call  ms/call  name
19.4    3.28    11962  0.27    _compare_strings
15.6    5.92    32731  0.08    _strlen
12.6    8.06    4579   0.47    _doprnt
10.5    9.84
9.9     11.52   6849   0.25    _get_field
5.3     12.42   762    1.18    _fgets
4.7     13.22   19715  0.04    _strcmp
4.0     13.89   5329   0.13    _malloc
3.4     14.46   11152  0.05    _insert_index_entry
3.1     14.99   11152  0.05    _compare_entry
2.5     15.41   1289   0.33    lmodt
0.9     15.57   761    0.21    _get_index_terms
0.9     15.73   3805   0.04    _strcpy
0.8     15.87   6849   0.02    _skip_space
0.7     15.99   13     9.23    _read
0.7     16.11   1289   0.09    ldivt
0.6     16.21   1405   0.07    _print_index
.
.
everything else is insignificant
.
.
```

This display points out that most of the program's running time is spent in the routine that compares character strings to establish the correct place for the index entries, and that after that, the majority of the time is spent in the `_strlen` library routine — to find the length of a character string. If we wish to make any appreciable improvements to the program we must concentrate our efforts on the `compare_strings` function.

Interpreting Profile Display

Let's interpret the results of the profiling run though. The results appear under these column headings:

```
%time  cumsecs  #call  ms/call  name
```

Here's what the columns mean:

- `%time` Percentage of the total run time of the program, that was consumed by this routine.
- `cumsecs` A running sum of the number of seconds accounted for by this function and those listed above it. This information isn't really worth much — the important data comes from the percentage of total time and from the time consumed per call.
- `#call` The number of times this routine was called.
- `ms/call` How many milliseconds this routine consumed each time it was called.

name The name of the routine.

Now what advice can we derive from the profile data? Notice that the `compare_strings` function consumes nearly 20% of the total time. To improve the run time of `index.assist` we must either improve the algorithm that `compare_strings` uses, or we must cut down the number of calls. Not obvious from the flat profile is the information that `compare_strings` is heavily recursive — we get that fact from using the call graph profile described below. In this particular case, improving the algorithm also implies reducing the number of calls.

5.3. `gprof` — Generate Call Graph Profile of Program

While the flat profile described in the last section can provide valuable data for performance improvements, sometimes the data obtained is not sufficient to point out exactly where the improvements can be made. A more detailed analysis can be obtained by using the *call graph* profile that displays a list of which modules are called by other modules, and which modules call other modules. Sometimes, removing calls altogether can result in performance improvements.

Compiling with the `-pg` Option

Using the same `index.assist` program as an example, let's make the program compiled for call-graph profiling. To compile a program for call-graph profiling, you use the `-pg` option to the C compiler:

```
tutorial% make CFLAGS=-pg
.
.
messages from the make command
.
.
tutorial%
```

Now we can run the `index.assist` program as before. When a program is call-graph profiled, the results appear in a file called `gmon.out` at the end of the run. You then use the `gprof` command to interpret the results of the profile:

```
tutorial% index.assist < index.entries > /dev/null
tutorial% gprof index.assist
.
.
voluminous output from the gprof command
.
.
```

Output from `gprof`

The output from `gprof` is really voluminous — it's usually intended that you take the summaries away and read them later. The output from `gprof` consists of two major items:

- The 'flat' profile. This is similar to the summary that the `prof` command supplies. `gprof` gives you slightly more information. The output from `gprof` contains an explanation of what the various parts of the summary

mean, so you don't need to go look the things up in a manual.

- The full call-graph profile. There are some fragments of the output from the profiling run just below with some examples of how to interpret them.

The output from `gprof` contains an explanation of what the various parts of the summary mean, so you don't need to go look the things up in a manual.

Interpreting Call Graph

Here is a fragment of the output from the `gprof` summary. Most of the output has been deleted from before and after the fragment. One thing that `gprof` does tell you is the granularity of the sampling:

granularity: each sample hit covers 4 byte(s) for 0.14% of 14.74 seconds

Then comes part of the call-graph profile itself:

index	%time	self	descendents	called/total	parents	
				called+self called/total	name	index children

[2]	98.2	0.00	14.47	1/1	start	[1]
		0.00	14.47	1	_main	[2]
		0.59	5.70	760/760	_insert_index_entry	[3]
		0.02	3.16	1/1	_print_index	[6]
		0.20	1.91	761/761	_get_index_terms	[11]
		0.94	0.06	762/762	_fgets	[13]
		0.06	0.62	761/761	_get_page_number	[18]
		0.10	0.46	761/761	_get_page_type	[22]
		0.09	0.23	761/761	_skip_start	[24]
		0.04	0.23	761/761	_get_index_type	[26]
		0.07	0.00	761/820	_insert_page_entry	[34]

[3]	42.6	0.59	5.70	10392	_insert_index_entry	[3]
		0.59	5.70	760/760	_main	[2]
		0.53	5.13	760+10392	_insert_index_entry	[3]
		0.02	0.01	11152/11152	_compare_entry	[4]
		0.02	0.01	59/112	_free	[38]
		0.00	0.00	59/820	_insert_page_entry	[34]
				10392	_insert_index_entry	[3]

Noting that there are 761 lines of data in the input file to the `index.assist` program, here are some of the things we can determine from the call graph:

- `fgets` is called 762 times — one more than the number of lines in the input file. The last call to `fgets` returns an end-of-file.
- The `insert_index_entry` function is called 760 times from `main` — one less times than the number of lines. Why is this? The first index entry is inserted 'manually' in the `main` function when there are no previous

index entries to insert.

- Note that in addition to the 760 times that `insert_index_entry` is called from `main`, `insert_index_entry` also calls *itself* the grand total of 10392 times — `insert_index_entry` is heavily recursive. Index entries appear in the input file in unsorted order and are sorted on the fly by inserting them into a binary tree.
- Note also that `compare_entry` (which is called from `insert_index_entry`) is called 11152 times, which is equal to $760+10392$ times, so there is one call of `compare_entry` for every time that `insert_index_entry` is called. This is as it should be. If there was a discrepancy in the number of calls, we might suspect some problem in the program's logic.
- Notice the number of calls to the `insert_page_entry` and `free` functions — `insert_page_entry` is called 820 times in total: 761 times from `main` while the program is building index nodes, and then `insert_page_entry` is called 59 times from `insert_index_entry`. This indicates that there are 59 index entries that are duplicated, so their page number entries are linked into a chain with the index nodes. The duplicate index entries are then freed, hence the 59 calls to `free`.

5.4. `tcov` — Statement-level Analysis of Program

After a certain level of performance enhancements have been made, the profile data obtained from a program starts to look 'flat' and the granularity of the data collection makes further improvements difficult. At this point, you can use a tool that performs statement-by-statement analysis on a program, showing which statements are executed and how many times. This facility is called *code coverage*.

Code coverage can also be valuable in identifying areas of 'dead' code — areas of code that never get executed. Code coverage can also point out areas of code that are not being tested.

Compiling with the `-a` Option

Using the same `index.assist` program as an example, let's make the program compiled for code coverage. To compile a program for code coverage, you use the `-a` option to the C compiler:

```
tutorial% make CFLAGS=-a
.
.
  messages from the make command
.
.
tutorial%
```

For every *thing.c* file you compile with the `-a` option, the C compiler generates a *thing.d* file — these are used by the code coverage program later in the analysis.

Using `tcov`

Now we can run the `index.assist` program as before. After a program has been run, you can then run `tcov` to get the summaries of execution counts for each statement in the program:

```
tutorial% index.assist < index.entries > /dev/null
tutorial% tcov *.c
```

Now, for every *thing.c* file you specify, `tcov` uses the *thing.d* file and generates a *thing.tcov* file containing an annotated listing of your code. The listing shows the number of times each source statement was executed. At the end of each *thing.tcov* file there is a short summary.

Here is a small fragment of the C code from one of the modules of `index.assist` — the module in question is the `insert_index_entry` function that's called so recursively:

```
struct index_entry *
insert_index_entry(node, entry)
11152 -> struct index_entry *node;
        struct index_entry *entry;
        {

        int result;
        int level;

        result = compare_entry(node, entry);

        if (result == 0) { /* exact match */
            /* Place the page entry for the duplicate */
            /* into the list of pages for this node */
59 ->    insert_page_entry(node, entry->page_entry);
        free(entry);
        return(node);
        }

11093 -> if (result > 0) /* node greater than new entry -- */
        /* move to lesser nodes */
3956 ->    if (node->lesser != NULL)
3626 ->        insert_index_entry(node->lesser, entry);
        else {
330 ->        node->lesser = entry;
        return (node->lesser);
        }
        else /* node less than new entry -- */
            /* move to greater nodes */
7137 ->    if (node->greater != NULL)
6766 ->        insert_index_entry(node->greater, entry);
        else {
371 ->        node->greater = entry;
        return (node->greater);
        }
    }
}
```

Notice that the `insert_index_entry` function is indeed called 11152 times as we determined in the output from `gprof`. The numbers to the side of the C code show how many times each statement was executed.

`tcov` Summary

Here is the summary that `tcov` placed at the end of `build.index.tcov`:

Top 10 Blocks

Line	Count
240	21563
241	21563
245	21563
251	21563
250	21400
244	21299
255	20612
257	16805
123	12021
124	11962

77	Basic blocks in this file
55	Basic blocks executed
71.43	Percent of the file executed
439144	Total basic block executions
5703.17	Average executions per basic block

1

m4 — A Macro Processor

m4 — A Macro Processor	107
6.1. Using the m4 Command	108
6.2. Defining Macros	108
6.3. Quoting and Comments	109
6.4. Macros with Arguments	110
6.5. Arithmetic Built-ins	111
6.6. File Manipulation	112
6.7. Running System Commands	113
6.8. Conditionals	113
6.9. String Manipulation	114
6.10. Printing	115
6.11. Summary of Built-in m4 Macros	115

m4 — A Macro Processor

m4 is a macro processor whose primary use has been as a front end for Ratfor for those cases where parameterless macros are not adequately powerful. It has also been used for languages as disparate as C and COBOL. m4 is particularly suited for higher-level languages like FORTRAN, PL/I and C since macros are specified in a functional notation.

m4 provides features seldom found even in much larger macro processors, including

- arguments
- condition testing
- arithmetic capabilities
- string and substring functions
- file manipulation

A macro processor is a useful way to enhance a programming language, to make it more palatable or more readable, or to tailor it to a particular application. The `#define` statement in C and the analogous `define` in Ratfor are examples of the basic facility provided by any macro processor — replacement of text by other text.

The basic operation of m4 is to copy its input to its output. As the input is read, however, each alphanumeric “token” (that is, string of letters and digits) is checked. If it is the name of a macro, then the name of the macro is replaced by its defining text, and the resulting string is pushed back onto the input to be rescanned. Macros may be called with arguments, in which case the arguments are collected and substituted into the right places in the defining text before it is rescanned.

m4 provides a collection of about twenty built-in macros which perform various useful operations; in addition, the user can define new macros. Built-in macros and user-defined macros work exactly the same way, except that some of the built-in macros have side effects on the state of the process.

6.1. Using the m4 Command

The basic m4 command line looks like this:

```
tutorial% m4 [ filename . . . ]
```

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 to the standard output, which may be captured for subsequent processing by redirecting the standard output:

```
tutorial% m4 [ filename . . . ] > outputfile
```

6.2. Defining Macros

The primary built-in function of m4 is `define`, which is used to define new macros. The input

```
define(name, stuff)
```

defines the string *name* as *stuff*. All subsequent occurrences of *name* will be replaced by *stuff*, unless *name* is redefined, or until *name* is undefined. *name* must be alphanumeric and must begin with a letter (the underscore `_` counts as a letter). *stuff* is any text that contains balanced parentheses; it may stretch over multiple lines.

Thus, as a typical example,

```
define(N, 100)
. . .
if (i > N)
```

defines *N* to be 100, and uses this "symbolic constant" in a later *if* statement.

The left parenthesis must immediately follow the word `define`, to signal that `define` has arguments. If a macro or built-in name is not followed immediately by '(', it is assumed to have no arguments. This is the situation for *N* above; it is actually a macro with no arguments, and thus when it is used there need be no parenthesis following it.

m4 divides its input into tokens, so a macro name is only recognized as such if it appears surrounded by non-alphanumerics. For example, in

```
define(N, 100)
. . .
if (NNN > 100)
```

the variable *NNN* is absolutely unrelated to the defined macro *N*, even though it contains a lot of *N*'s.

Things may be defined in terms of other things. For example,

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

defines both *M* and *N* to be 100.

What happens if N is redefined? Or, to say it another way, is M defined as N or as 100? In m4, the latter is true — M is 100, so that changing N does not change M .

This behavior arises because m4 expands macro names into their defining text as soon as it possibly can. Here, that means that when the string N is seen while the arguments of `define` are being collected, it is immediately replaced by 100; it's just as if you had said

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

in the first place.

If this isn't what you really want, there are two ways out of it. The first, which is specific to this situation, is to interchange the order of the definitions:

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

Now M is defined to be the string N , so when you ask for M later, you'll always get the value of N at that time (because the M will be replaced by N which will be replaced in turn by its value).

6.3. Quoting and Comments

The more general solution is to delay the expansion of the arguments of `define` by *quoting* them. Any text surrounded by the single quotes ``` and `'` is not expanded immediately, but has the quotes stripped off. If you say

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

the quotes around the N are stripped off as the argument is being collected, but they have served their purpose, and M is defined as the string N , not 100. The general rule is that m4 always strips off one level of single quotes whenever it evaluates something. This is true even outside of macros. If you want the word `define` to appear in the output, you have to quote it in the input, as in

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

As another instance of the same thing, which is a bit more surprising, consider redefining N :

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

Perhaps regrettably, the N in the second definition is evaluated as soon as it's seen; that is, it is replaced by 100, so it's as if you had written

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

This statement is ignored by m4, since you can only define things that look like names, but it obviously doesn't have the effect you wanted. To really redefine N , you must delay the evaluation by quoting:

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

If ``` and `'` are not convenient for some reason, the quote characters can be changed with the built-in `changequote`:

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

makes the new quote characters the left and right brackets. You can restore the original characters with just

```
changequote
```

There are two additional built-ins related to `define`. `undefine` removes the definition of some macro or built-in:

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

removes the definition of *N*. (Why are the quotes absolutely necessary?) Built-ins can be removed with `undefine`, as in

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

but once you remove one, you can never get it back.

The built-in `ifdef` provides a way to determine if a macro is currently defined. In particular, `m4` pre-defines the name *unix*.

`ifdef` actually permits three arguments; if the name is undefined, the value of `ifdef` is then the third argument, as in

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

Don't forget the quotes around the argument.

Comments in `m4` are introduced by the `#` (sharp) character. All text from the `#` to the end of the line is taken as a comment and otherwise ignored.

6.4. Macros with Arguments

So far we have discussed the simplest form of macro processing — replacing one string by another (fixed) string. User-defined macros may 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 macro *bump*, defined as

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

generates code to increment its argument by 1:

```
bump(x)
```

evaluates to

```
x = x + 1
```

A macro can have as many arguments as you want, but only the first nine are accessible, through $\$1$ to $\$9$. The macro name itself is $\$0$, although that is less commonly used. Arguments that are not supplied are replaced by null strings, so

we can define a macro *cat* which 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
```

\$4 through *\$9* are null, since no corresponding arguments were 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, but commas can be nested inside parentheses. That is, in

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

there are only two arguments; the second is literally *(b,c)*. And of course a bare comma or parenthesis can be inserted by quoting it.

6.5. Arithmetic Built-ins

m4 provides two built-in functions for doing arithmetic on integers (only). The simplest is *incr*, which increments its numeric argument by 1. Thus to handle the common programming situation where you want a variable to be defined as “one more than *N*”, write

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

which defines *N1* as one more than the current value of *N*.

The more general mechanism for arithmetic is a built-in called *eval*, which is capable of arbitrary arithmetic on integers. *eval* provides the operators (in decreasing order of precedence)

Table 6-1 Operators to the eval built in in m4

<i>Operator</i>	<i>Meaning</i>
unary + and -	add and subtract
** or ^	exponentiation
* / %	multiply, divide, and modulus
+ -	binary add and subtract
== != < <= > >=	equal, not equal, less than, less than or equal, greater than, greater than or equal
!	logical not
& or &&	logical and)
or	(logical or)

Parentheses may 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.

As a simple example, suppose we want M to be 2^{**N+1} . Then

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

As a matter of principle, it is advisable to quote the defining text for a macro unless it is very simple indeed (say just a number); it usually gives the result you want, and is a good habit to get into.

6.6. File Manipulation

You can include a new file in the input at any time by the built-in function `include`:

```
include(filename)
```

inserts the contents of *filename* in place of the `include` command. The contents of the file is often a set of definitions. The value of `include` (that is, its replacement text) is the contents of the file; this can be captured in definitions, etc.

It is a fatal error if the file named in `include` cannot be accessed. To get some control over this, the alternate form `sinclude` can be used; `sinclude` ("silent include") says nothing and continues if it can't access the file.

It is also possible to divert the output of `m4` to temporary files during processing, and output the collected material upon command. `m4` maintains nine of these diversions, numbered 1 through 9. If you say

```
divert(n)
```

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

Diverted text is normally output all at once at the end of processing, with the diversions output in numeric order. It is possible, however, to bring back diversions at any time, that is, to append them to the current diversion.

```
undivert
```

brings back all diversions in numeric order, and `undivert` with arguments brings back the selected diversions in the order given. The act of undiverting discards the diverted stuff, as does diverting into a diversion whose number is not between 0 and 9 inclusive.

The value of `undivert` is *not* the diverted stuff. Furthermore, the diverted material is *not* rescanned for macros.

The built-in `divnum` returns the number of the currently active diversion. This is zero during normal processing.

6.7. Running System Commands

You can run any UNIX† program with the `syscmd` built-in. For example,

```
syscmd(date)
```

runs the *date* command. Normally `syscmd` would be used to create a file for a subsequent `include`.

To facilitate making unique file names, the built-in `maketemp` is provided, with specifications identical to the system function `mktemp`: a string of `XXXXXX` in the argument is replaced by the process id of the current process.

6.8. Conditionals

There is a built-in called `ifelse` which enables you to perform arbitrary conditional testing. In its simplest form,

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

compares the two strings *a* and *b*. If these are identical, `ifelse` returns the string *c*; otherwise it returns *d*. Thus we might define a macro called *compare* which compares two strings and returns “yes” or “no” according to whether they are the same or different.

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

Note the quotes, which prevent too-early evaluation of `ifelse`.

If the fourth argument is missing, it is treated as empty.

`ifelse` can actually have any number of arguments, and thus provides a limited form of multi-way decision capability. In the input

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

if the string *a* matches 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, the result is null, so

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

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

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

6.9. String Manipulation

The built-in `len` returns the length of the string that makes up its argument. Thus

```
len(abcdef)
```

is 6, and `len((a,b))` is 5.

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

`substr(s, i, n)` returns the substring of *s* that starts at the *i*th position (origin zero), and is *n* characters long. If *n* is omitted, the rest of the string is returned, so

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

evaluates to

```
ow is the time
```

If either *i* or *n* is out of range, various sensible things happen.

`index(s1, s2)` returns the index (position) in *s1* where the string *s2* occurs, or `-1` if it doesn't occur. As with `substr`, the origin for strings is 0.

The built-in `translit` performs character transliteration.

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

modifies *s* by replacing any character found in *f* by the corresponding character in *t*. That is,

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

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

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

deletes vowels from *s*.

There is also a built-in called `dn1` which deletes all characters that follow it up to and including the next newline; it is useful mainly for throwing away empty lines that otherwise tend to clutter up `m4` output. For example, if you say

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

the newline at the end of each line is not part of the definition, so it is copied into the output, where it may not be wanted. If you add `dn1` to each of these lines, the newlines will disappear.

Another way to achieve this, due to J. E. Weythman, is

```

divert (-1)
  define(... )
  ...
divert

```

6.10. Printing

The built-in `errprint` writes its arguments to the standard error file. Thus you can say

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

`dumpdef` is a debugging aid which dumps the current definitions of defined terms. If there are no arguments, you get everything; otherwise you get the ones you name as arguments. Don't forget to quote the names!

6.11. Summary of Built-in m4 Macros

Table 6-2 *Summary of Built-in m4 Macros*

<i>Built In</i>	<i>Description</i>
<code>changequote(L, R)</code>	Change left quote to L, right quote to R
<code>define(name, replacement)</code>	define <i>name</i> as <i>replacement</i>
<code>divert(number)</code>	Divert output to stream <i>number</i>
<code>divnum</code>	Return number of currently active diversions
<code>dnl</code>	Delete up to and including new-line
<code>dumpdef('name', 'name', ...)</code>	Dump specified definitions
<code>errprint(s, s, ...)</code>	Write arguments <i>s</i> to standard error
<code>eval(numeric expression)</code>	Evaluate <i>numeric expression</i>
<code>ifdef('name', true string, false string)</code>	Return <i>true string</i> if <i>name</i> is defined, <i>false string</i> if <i>name</i> is not defined
<code>ifelse(a, b, c, d)</code>	If <i>a</i> and <i>b</i> are equal, return <i>c</i> , else return <i>d</i>
<code>include(file)</code>	Include contents of <i>file</i>
<code>incr(number)</code>	Increment <i>number</i> by 1
<code>index(s1, s2)</code>	Return position in <i>s1</i> where <i>s2</i> occurs, or -1 if no occurrence
<code>len(string)</code>	Return length of <i>string</i>

Table 6-2 Summary of Built-in m4 Macros—Continued

<i>Built In</i>	<i>Description</i>
<code>maketemp(. . .XXXXX. . .)</code>	Make a temporary file
<code>sinclude (file)</code>	Include contents of <i>file</i> — ignored and continue if <i>file</i> not found.
<code>substr (string, position, number)</code>	Return substring of <i>string</i> starting at <i>position</i> and <i>number</i> characters long
<code>syscmd (command)</code>	Run <i>command</i> in the system
<code>translit (string, from, to)</code>	Transliterate characters in <i>string</i> from the set specified by <i>from</i> to the set specified by <i>to</i>
<code>undefine ('name')</code>	Remove <i>name</i> from the list of definitions
<code>undivert (number, number, . . .)</code>	Append diversion <i>number</i> to the current diversion

Lex — A Lexical Analyzer Generator

Lex — A Lexical Analyzer Generator	119
7.1. Lex Source	122
7.2. Lex Regular Expressions	123
7.3. Lex Actions	126
7.4. Ambiguous Source Rules	129
7.5. Lex Source Definitions	131
7.6. Using <code>lex</code>	132
7.7. Lex and Yacc	133
7.8. Examples	133
7.9. Left Context-Sensitivity	136
7.10. Character Set	138
7.11. Summary of Source Format	138
7.12. Caveats and Bugs	140

Lex — A Lexical Analyzer Generator

`lex` is a program generator designed for lexical processing of character input streams. `lex` accepts a high-level, problem-oriented specification for character string matching, and produces a program in a general-purpose language which recognizes regular expressions. The regular expressions are specified by the programmer in the source specifications given to `lex`. The `lex` written code recognizes these expressions in an input stream and partitions the input stream into strings matching the expressions. At the boundaries between strings, program sections provided by the programmer are executed. The `lex` source file associates the regular expressions and the program fragments. As each expression appears in the input to the program written by `lex`, the corresponding fragment is executed.

The programmer supplies the additional code beyond expression matching needed to complete his tasks, possibly including code written by other generators. The program that recognizes the expressions is generated in the general-purpose programming language employed for the programmer's program fragments. Thus, a high-level expression language is provided to write the string expressions to be matched while the programmer's freedom to write actions is unimpaired. This avoids forcing the programmer who wishes to use a string manipulation language for input analysis to write processing programs in the same and often inappropriate string handling language.

`lex` source is a table of regular expressions and corresponding program fragments. The table is translated to a program which reads an input stream, copying it to an output stream and partitioning the input into strings which match the given expressions. As each such string is recognized the corresponding program fragment is executed. The recognition of the expressions is performed by a deterministic finite automaton generated by `lex`. The program fragments written by the programmer are executed in the order in which the corresponding regular expressions occur in the input stream.

The lexical analysis programs written with `lex` accept ambiguous specifications and choose the longest match possible at each input point. If necessary, substantial lookahead is performed on the input, but the input stream is then backed up to the end of the current partition, so that the programmer has general freedom to manipulate it.

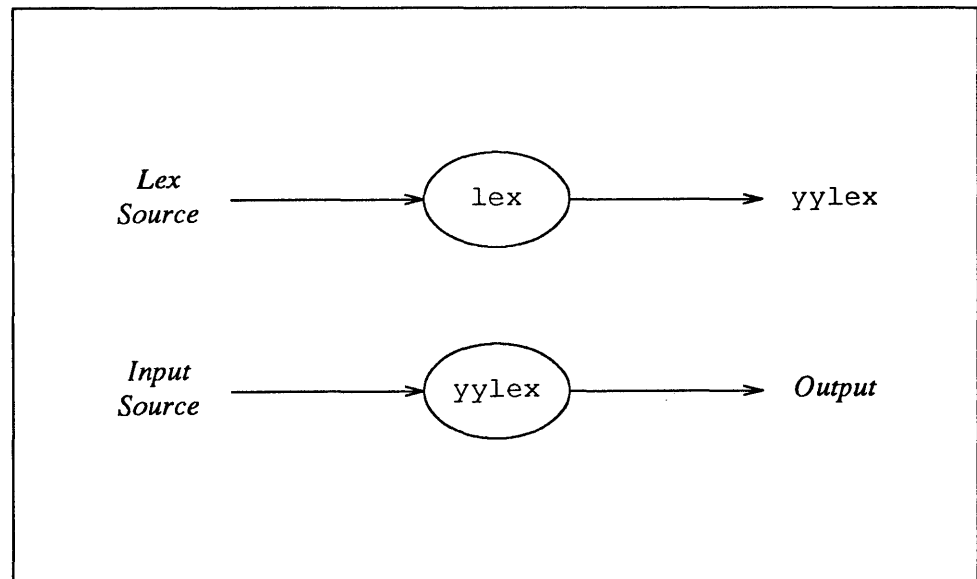
`lex` can generate analyzers in either C or Ratfor, a language which can be translated automatically to portable FORTRAN. `lex` is designed to simplify

interfacing with `yacc`, which is described in the next chapter.

`lex` is not a complete language, but rather a generator representing a new language feature which can be added to different programming languages, called 'host languages.' 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 also for the program fragments added by the programmer. Compatible run-time libraries for the different host languages are also provided. This makes `lex` adaptable to different environments and different programmer. Each application may be directed to the combination of hardware and host language appropriate to the task, the programmer's background, and the properties of local implementations.

`lex` turns the programmer's expressions and actions (called *source* in this document) into the host general-purpose language; the generated program is named `yylex`. The `yylex` program recognizes expressions in a stream (called *input* in this document) and performs the specified actions for each expression as it is detected — see Figure 7-1 below.

Figure 7-1 *An overview of Lex*



For a trivial example, consider a program to delete from the input all blanks or tabs at the ends of lines.

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

is all that is required. The program contains a `%%` delimiter to mark the beginning of the rules, and one rule. This rule contains a regular expression which matches one or more instances of the characters blank or tab (written `\t` for visibility, in accordance with the C convention) just prior to the end of a line. The brackets indicate the character class made of blank and tab; the `+` indicates 'one or more ...'; and the `$` indicates 'end-of-line'. No action is specified, so the program generated by `lex` (`yylex`) ignores these characters. Everything else is

copied to the output stream. To change any remaining string of blanks or tabs to a single blank, add another rule:

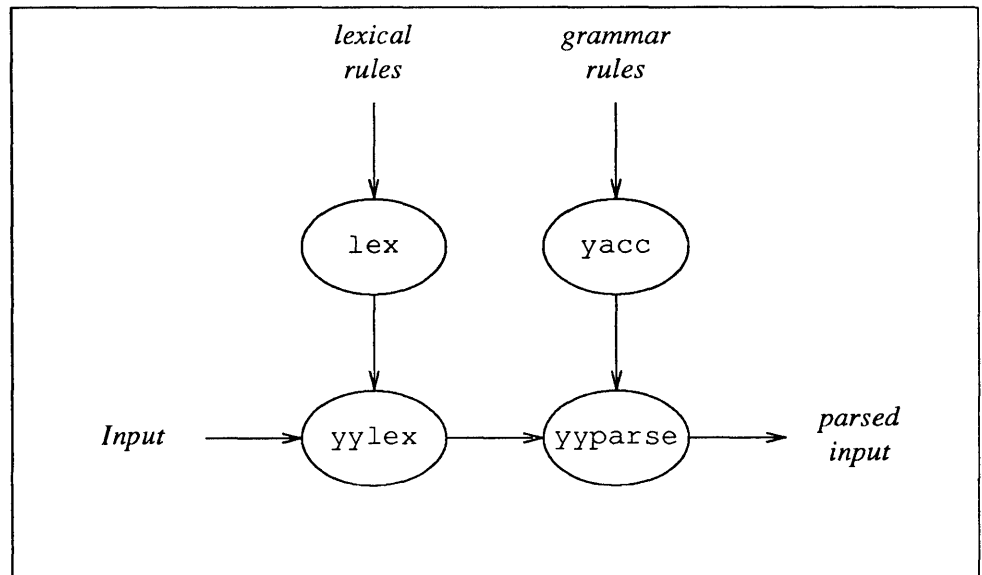
```
%%
[ \t]+$ ;
[ \t]+ printf(" ");
```

The finite automaton generated for this source scans for both rules at once, observing at the termination of the string of blanks or tabs whether or not there is a newline character, and executing the desired rule action. The first rule matches all strings of blanks or tabs at the ends of lines, and the second rule all remaining strings of blanks or tabs.

`lex` can also be used with a parser generator to perform the lexical analysis phase.

`lex` can be used alone for simple transformations, or for analysis and statistics gathering on a lexical level. `lex` can also be used with a parser generator to perform the lexical analysis phase; it is particularly easy to interface `lex` and `yacc`[3]. `lex` programs recognize only regular expressions; `yacc` writes parsers that accept a large class of context-free grammars, but require 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, and the parser generator assigns structure to the resulting pieces. The flow of control in such a case (which might be the first half of a compiler, for example) is shown in Figure 7-2. Additional programs, written by other generators or by hand, can be added easily to programs written by `lex`.

Figure 7-2 *Lex with Yacc*



`yacc` programmers will realize that the name `yylex` is what `yacc` expects its lexical analyzer to be named, so that the use of this name by `lex` simplifies interfacing.

`lex` generates a deterministic finite automaton from the regular expressions in the source [4]. The automaton is interpreted, rather than compiled, in order to save space. The result is still a fast analyzer. In particular, the time taken by a `lex` program to recognize and partition an input stream is proportional to the length of the input. The number of `lex` rules or the complexity of the rules is not important in determining speed, unless rules which include forward context require a significant amount of rescanning. What does increase with the number and complexity of rules is the size of the finite automaton, and therefore the size of the program generated by `lex`.

In the program written by `lex`, the programmer's fragments (representing 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. Opportunity is provided for the programmer to insert either declarations or additional statements in the routine containing the actions, or to add subroutines outside this action routine.

`lex` is not limited to source which can be interpreted on the basis of one character lookahead. 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 leave the input pointer just before "`cd...`" Such backup is more costly than processing simpler languages.

7.1. Lex Source

The general format of `lex` source is:

```
{definitions}
%%
{rules}
%%
{programmer subroutines}
```

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

```
%%
```

(no definitions, no rules) which translates into a program which copies the input to the output unchanged.

In the outline of `lex` programs shown above, the *rules* represent the programmer's control decisions; they are a table, in which the left column contains *regular expressions* (see section 7.2) and the right column contains *actions*, program fragments to be executed when the expressions

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

to look for the string `integer` in the input stream and print the message 'found keyword INT' whenever it appears. In this example the host procedural language is C and the C 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 merely a single C expression, it can just be given on the right side of the line; if it is compound, or takes more than a line, it should be enclosed in braces. As a slightly more useful example, suppose it is desired to change a number of words from British to American spelling. `lex` rules such as

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

would be a start. These rules are not quite enough, since the word *petroleum* would become *gaseum*; a way of dealing with this is described later.

7.2. Lex Regular Expressions

The definitions of regular expressions are very similar to those in the UNIX editors *ex*(1) and *vi*(1)[5]. A regular expression specifies a set of strings to be matched. It contains text characters (which match the corresponding characters in the strings being compared) and operator characters (which specify repetitions, choices, and other features). The letters of the alphabet and the digits are always text characters; thus the regular expression

```
integer
```

matches the string *integer* wherever it appears and the expression

```
a57D
```

looks for the string *a57D*.

Operators

The operator characters are

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

and if they are to be used as text characters, an escape must be used. The quotation mark operator ("") indicates that whatever is contained between a pair of quotes is to be taken as text characters. Thus

```
xyz"++"
```

matches the string *xyz++* when it appears. Note that a part of a string may be quoted. It is harmless but unnecessary to quote an ordinary text character; the expression

```
"xyz++"
```

is the same as the one above. Thus by quoting every non-alphanumeric character being used as a text character, the programmer can avoid remembering the list above of current operator characters, and is safe should further extensions to `lex` lengthen the list.

An operator character may also be turned into a text character by preceding it with `\` as in

```
xyz\+\+
```

which is another, less readable, equivalent of the above expressions. Another use of the quoting mechanism is to get a blank into an expression; normally, as explained above, blanks or tabs end a rule. Any blank character not contained

within `[]` (see below) must be quoted. Several normal C escapes with `\` are recognized: `\n` is newline, `\t` is tab, and `\b` is backspace. To enter `\` itself, use `\\`. Since newline is illegal in an expression, `\n` must be used; it is not required to escape tab and backspace. Every character but blank, tab, newline and the list above is always a text character.

Character classes

Classes of characters can be specified using the operator pair `[]`. The construction `[abc]` matches a single character, which may be *a*, *b*, or *c*. Within square brackets, most operator meanings are ignored. Only three characters are special: these are `\`, `-`, and `^`. The `-` character indicates ranges. For example,

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

indicates the character class containing all the lower case letters, the digits, the angle brackets, and underline. Ranges may be given in either order. Using `-` between any pair of characters which are not both upper case letters, both lower case letters, or both digits is implementation-dependent and generates a warning message. For example, `[0-z]` in ASCII is many more characters than it is in EBCDIC. If it is desired to include the character `-` in a character class, it should be first or last, thus:

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

matches all the digits and the two signs.

In character classes, the `^` operator must appear as the first character after the left bracket; it indicates that the resulting string is to be complemented with respect to the system's character set. Thus

```
[^abc]
```

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

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

is any character which is not a letter. The `\` character provides the usual escapes within character class brackets.

Arbitrary character

To match almost any character, the operator character

(period) is the class of all characters except newline. Escaping into octal is possible although non-portable:

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

matches all printable characters in the ASCII character set, from octal 40 (blank) to octal 176 (tilde).

Optional expressions

The operator `?` indicates 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 `+`.

`a*`

is any number of consecutive *a* characters, including zero; while

`a+`

is one or more instances of *a*. For example,

`[a-z]+`

is all strings of lower case letters. And

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

indicates all alphanumeric strings with a leading alphabetic character. This is a typical expression for recognizing identifiers in computer languages.

Alternation and Grouping

The operator `|` indicates alternation:

`(ab | cd)`

matches either *ab* or *cd*. Note that parentheses are used for grouping, although they are not necessary on the outside level;

`ab | cd`

would have sufficed. Parentheses can be used for more complex expressions:

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

matches such strings as *abefef*, *efefef*, *cdef*, or *cddd*; but not *abc*, *abcd*, or *abcdef*.

Context sensitivity

`lex` recognizes a small amount of surrounding context. The two simplest operators for this are `^` and `$`. If the first character of an expression is `^`, the expression is only be matched at the beginning of a line This can never conflict with the other meaning of `^`, complementation of character classes, since that only applies within the `[]` operators. If the very last character is `$`, the expression is only be matched at the end of a line (when immediately followed by newline). The latter operator is a special case of the `/` operator character, which indicates trailing context. The expression

`ab/cd`

matches the string *ab*, but only if it is followed by *cd*. Thus

`ab$`

is the same as

`ab/\n`

Left context is handled in `lex` by *start conditions* as explained in section 7.9 — *Left Context-Sensitivity*. If a rule is only to be executed when the `lex` automaton interpreter is in start condition *x*, the rule should be prefixed by

`<x>`

using the angle bracket operator characters. If we considered ‘being at the

beginning of a line' to be start condition *ONE*, then the `^` operator would be equivalent to

```
<ONE>
```

Start conditions are explained more fully below.

Repetitions and Definitions

The operators `{ }` specify either repetitions (if they enclose numbers) or 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. In contrast,

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

looks for 1 to 5 occurrences of *a*.

Finally, initial `%` is special, being the separator for `lex` source segments.

7.3. Lex Actions

When an expression written as above is matched, `lex` executes the corresponding action. This section describes some features of `lex` which aid in writing actions. Note that there is a default action, which consists of copying the input to the output. This is performed on all strings not otherwise matched. Thus the `lex` programmer who wishes to absorb the entire input, without producing any output, must provide rules to match everything. When `lex` is being used with `yacc`, this is the normal situation. One may consider that actions are what is done instead of copying the input to the output; thus, in general, a rule which merely copies can be omitted. Also, a character combination which is omitted from the rules and which appears as input is likely to be printed on the output, thus calling attention to the gap in the rules.

One of the simplest things that can be done is to ignore the input. Specifying a C null statement, `;` as an action does this. A frequent rule is

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

which ignores the three spacing characters (blank, tab, and newline).

Another easy way to avoid writing actions is the action character `|`, which indicates that the action to be used for this rule is the action given for the next rule. The previous example could also have been written

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

with the same result. The quotes around `\n` and `\t` are not required.

Actual

In more complex actions, the programmer often wants to know the actual text that matched some expression like `[a-z]+`. `lex` leaves this text in an external character array named `ytext`. Thus, to print the name found, a rule like

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

prints the string in `yytext`. The C function `printf` accepts a format argument and data to be printed; in this case, the format is 'print string' (`%` indicating data conversion, and `s` indicating string type), and the data are the characters in `yytext`. So this just places the matched string on the output. This action is so common that it may be written as ECHO:

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

is the same as the above. Since the default action is just to print the characters found, one might ask why give a rule, like this one, which merely specifies the default action? Such rules are often required to avoid matching some other rule which is not desired. For example, if there is a rule which matches *read* it normally matches the instances of *read* contained in *bread* or *readjust*; to avoid this, a rule of the form `[a-z]+` is needed. This is explained further below.

Length

Sometimes it is more convenient to know the end of what has been found; hence `lex` also provides a count `yylen` of the number of characters matched. To count both the number of words and the number of characters in words in the input, the programmer might write

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

which accumulates in `chars` the number of characters in the words recognized. The last character in the string matched can be accessed by

```
yytext[yylen-1]
```

yymore

Occasionally, a `lex` action may decide that a rule has not recognized the correct span of characters. Two routines are provided to aid with this situation. First, `yymore()` can be called to indicate that the next input expression recognized is to be tacked on to the end of this input. Normally, the next input string would overwrite the current entry in `yytext`. Second, `yyless(n)` may be called to indicate that not all the characters matched by the currently successful expression are wanted right now. The argument `n` indicates the number of characters to be retained in `yytext`. Further characters previously matched are returned to the input. This provides the same sort of lookahead offered by the `/` operator, but in a different form.

Example: Consider a language which defines a string as a set of characters between quotation (") marks, and provides that to include a " in a string it must be preceded by a \. The regular expression which matches that is somewhat confusing, so that it might be preferable to write

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

which, when faced with a string such as `"abc\def"` first matches the five characters `"abc\`; then the call to `yymore()` tacks the next part of the string, `"def,`

onto the end. Note that the final quote terminating the string should be picked up in the code labeled 'normal processing'.

The function `yyles()` might be used to reprocess text in various circumstances. Consider the problem of resolving (in old-style C) the ambiguity of `'=-a'`. Suppose it is desired to treat this as `'=- a'` but print a message. A rule might be

```

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

```

which prints a message, returns the letter after the operator to the input stream, and treats the operator as `'=-'`. Alternatively it might be desired to treat this as `'= -a'`. To do this, just return the minus sign as well as the letter to the input:

```

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

```

performs the other interpretation. Note that the expressions for the two cases might more easily be written

```

=-/[A-Za-z]

```

in the first case and

```

=/[A-Za-z]

```

in the second; no backup would be required in the rule action. It is not necessary to recognize the whole identifier to observe the ambiguity. The possibility of `'=-3'`, however, makes

```

=-/[^ \t\n]

```

a still better rule.

In addition to these routines, `lex` also permits access to the I/O routines it uses. They are:

- 1) `input()` which returns the next input character;
- 2) `output(c)` which writes the character `c` on the output; and
- 3) `unput(c)` pushes the character `c` back onto the input stream to be read later by `input()`.

By default these routines are provided as macro definitions, but the programmer can override them and supply private versions. These routines define the relationship between external files and internal characters, and must all be retained or modified consistently. They may be redefined, to transmit input or output to or from strange places, including other programs or internal memory; but the character set used must be consistent in all routines; a value of zero returned by `input` must mean end of file; and the relationship between `unput` and

input must be retained or the `lex` lookahead will not work. `lex` does not look ahead at all if it does not have to, but every rule ending in `+` `*` `?` or `$` or containing `/` implies lookahead. Lookahead is also necessary to match an expression that is a prefix of another expression. See section 7.10 for a discussion of the character set used by `lex`. The standard `lex` library imposes a 100-character limit on backup.

Another `lex` library routine that the programmer will sometimes want to redefine is `yywrap()` which is called whenever `lex` reaches an end-of-file. If `yywrap` returns a 1, `lex` continues with the normal wrapup on end of input. Sometimes, however, it is convenient to arrange for more input to arrive from a new source. In this case, the programmer should provide a `yywrap` which arranges for new input and returns 0. This instructs `lex` to continue processing. The default `yywrap` always returns 1.

This routine is also a convenient place to print tables, summaries, etc. at the end of a program. Note that it is not possible to write a normal rule which 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, since a value of 0 returned by `input` is taken to be end-of-file.

7.4. Ambiguous Source Rules

`lex` can handle ambiguous specifications. When more than one expression can match the current input, `lex` chooses as follows:

- 1) The longest match is preferred.
- 2) Among rules which matched the same number of characters, the rule given first is preferred.

Thus, suppose the rules

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

to be given in that order. If the input is *integers*, it is taken as an identifier, because `[a-z]+` matches 8 characters while `integer` matches only 7. If the input is *integer*, both rules match 7 characters, and the keyword rule is selected because it was given first. Anything shorter (for example, *int*) will not match the expression `integer` and so the identifier interpretation is used.

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

```
'.*'
```

might seem a good way of recognizing a string in single quotes. But it is an invitation for the program to read far ahead, looking for a distant single quote.

Presented with the input

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

the above expression matches

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

which is probably not what was wanted. A better rule is of the form

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

which, on the above input, stops after *'first'*. The consequences of errors like this are mitigated by the fact that the `.` operator does not match newline. Thus expressions like `.*` stop on the current line. Don't try to defeat this with expressions like `[.\\n]+` or equivalents; the `lex` generated program will try to read the entire input file, causing internal buffer overflows.

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. For example, suppose it is desired to count occurrences of both *she* and *he* in an input text. Some `lex` rules to do this might be

```
she s++;
he h++;
\\n |
. ;
```

where the last two rules ignore everything besides *he* and *she*. Remember that `.` does not include newline. Since *she* includes *he*, `lex` will normally *not* recognize the instances of *he* included in *she*, since once it has passed a *she* those characters are gone.

Sometimes the programmer would like to override this choice. The action `REJECT` means 'go do the next alternative.' It executes whatever rule was second choice after the current rule. The position of the input pointer is adjusted accordingly. Suppose the programmer really wants to count the included instances of *he*:

```
she {s++; REJECT;}
he {h++; REJECT;}
\\n |
. ;
```

these rules are one way of changing the previous example to do just that. After counting each expression, it is rejected; whenever appropriate, the other expression is then counted. In this example, of course, the programmer could note that *she* includes *he* but not vice versa, and omit the `REJECT` action on *he*; in other cases, however, it would not be possible *a priori* to tell which input characters were in both classes.

Consider the two rules

```
a[bc]+ { ... ; REJECT;}
a[cd]+ { ... ; REJECT;}

```

If the input is *ab*, only the first rule matches, and on *ad* only the second matches. The input string *accb* matches the first rule for four characters and then the second rule for three characters. In contrast, the input *accd* agrees with the second rule for four characters and the first rule for three.

In general, `REJECT` is useful whenever the purpose of `lex` is not to partition the input stream but to detect all examples of some items in the input, and the instances of these items may overlap or include each other. Suppose a digram table of the input is desired; 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, the appropriate source is

```
%%
[a-z][a-z] {digram[yytext[0]][yytext[1]]++; REJECT;}
\n ;
```

where the REJECT is necessary to pick up a letter pair beginning at every character, rather than at every other character.

7.5. Lex Source Definitions

Remember the format of the `lex` source:

```
{definitions}
%%
{rules}
%%
{programmer routines}
```

So far only the rules have been described. The programmer needs additional options, though, to define variables for use in his program and for use by `lex`. These can go either in the definitions section or in the rules section.

Remember that `lex` is turning the rules into a program. Any source not intercepted by `lex` is copied into the generated program. There are three classes of such things.

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

As a side effect of the above, lines which begin with a blank or tab, and which 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.

- 2) Anything included between lines containing only the delimiters `%{` and `%}` is copied out as above. 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.
- 3) Anything after the third `%%` delimiter, regardless of formats, etc., is copied out 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
```

and it associates the string given as a translation 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 then be invoked by the `{name}` syntax in

a rule. Using {D} for the digits and {E} for an exponent field, for example, might abbreviate rules to recognize numbers:

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

Note the first two rules for real numbers; both require a decimal point and contain an optional exponent field, but the first requires at least one digit before the decimal point and the second requires at least one digit after the decimal point. To correctly handle the problem posed by a FORTRAN expression such as 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.

The definitions section may also 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. These possibilities are discussed below under section 7.11 — *Summary of Source Format*.

7.6. Using `lex`

There are two steps in compiling a `lex` source program. First, the `lex` source must be turned into a generated program in the host general-purpose language. Then this program must be compiled and loaded, usually with a library of `lex` subroutines. The generated program is on a file named `lex.yy.c`. The I/O library is defined in terms of the C standard library in section 3 of the *UNIX Interface Reference Manual for the Sun Workstation*.

The `lex` library is accessed by the loader flag `-ll`. So an appropriate set of commands is

```
tutorial% lex source
tutorial% cc lex.yy.c -ll
tutorial%
```

The resulting program is placed on the usual file `a.out` for later execution. To use `lex` with `yacc` see below. Although the default `lex` I/O routines use the C standard library, the `lex` automata themselves do not do so; if private versions of `input`, `output`, and `unput` are given, the library can be avoided. `lex` has several options which are described in the `lex(1)` manual page.

7.7. Lex and Yacc

If you want to use `lex` with `yacc`, note that what `lex` writes is a program named `yylex()`, the name required by `yacc` for its analyzer. Normally, the default main program in the `lex` library calls this routine, but if `yacc` is loaded, and its main program is used, `yacc` calls `yylex()`.

In this case each `lex` rule should end with

```
return(token);
```

to return the appropriate token value.

An easy way to get access to `yacc`'s names for tokens 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 `yacc` input. Supposing the grammar to be named 'good' and the lexical rules to be named 'better' the UNIX command sequence can just be:

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

The `lex` and `yacc` programs can be generated in either order.

7.8. Examples

As a trivial problem, consider copying an input file while adding 3 to every non-negative number divisible by 7. Here is a suitable `lex` source program

```
%%
int k;
[0-9]+ {
    k = atoi(yytext);
    if (k%7 == 0)
        printf("%d", k+3);
    else
        printf("%d", k);
}
```

to do just that. The rule `[0-9]+` recognizes strings of digits; `atoi` converts the digits to binary and stores the result in `k`. The operator `%` (remainder) is used to check whether `k` is divisible by 7; if it is, it is incremented by 3 as it is written out. It may be objected that this program will alter such input items as `49.63` or `X7`. Furthermore, it increments the absolute value of all negative numbers divisible by 7. To avoid this, just add a few more rules after the active one, as here:

```
%%
    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;
```

Numerical strings containing a '.' or preceded by a letter are picked up by one of the last two rules, and not changed. The *if-else* has been replaced by a C conditional expression to save space; the form *a?b:c* means 'if *a* then *b* else *c*'.

For an example of statistics gathering, here is a program which constructs a histogram of the lengths of words, where a word is defined as a string of letters.

```
    int lengs[100];
%%
[a-z]+    lengs[yyval]++;
.    |
\n    ;
%%
l s.
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);
}
```

This program accumulates the histogram, while producing no output. At the end of the input it prints the table. The final statement *return(1)*; indicates that *lex* is to perform wrapup. If *yywrap* returns zero (false) it implies that further input is available and the program is to continue reading and processing. To provide a *yywrap* that never returns true causes an infinite loop.

As a larger example, here are some parts of a program written by N. L. Schryer to convert double-precision FORTRAN to single-precision FORTRAN. Because FORTRAN does not distinguish upper and lower case letters, this routine begins by defining a set of classes including both cases of each letter:

```
a    [aA]
b    [bB]
c    [cC]
...
z    [zZ]
```

An additional class recognizes white space:

```
W    [ \t]*
```

The first rule changes double precision to real, or DOUBLE

PRECISION to REAL.

```
{d}{o}{u}{b}{l}{e}{W}{p}{r}{e}{c}{i}{s}{i}{o}{n} {
    printf(yytext[0]=='d'? "real" : "REAL");
}
```

Care is taken throughout this program to preserve the case (upper or lower) of the original program. The conditional operator is used to select the proper form of the keyword. The next rule copies continuation card indications to avoid confusing them with constants:

```
^"      "[^ 0]    ECHO;
```

In the regular expression, the quotes surround the blanks. It is interpreted as 'beginning of line, then five blanks, then anything but blank or zero.' Note the two different meanings of `^`. There follow some rules to change double-precision constants to ordinary floating constants.

```
[0-9]+{W}{d}{W}[+-]?{W}[0-9]+      |
[0-9]+{W}"."{W}{d}{W}[+-]?{W}[0-9]+  |
"."{W}[0-9]+{W}{d}{W}[+-]?{W}[0-9]+  {
    /* convert constants */
    for(p=yytext; *p != 0; p++)
    {
        if (*p == 'd' || *p == 'D')
            *p+= 'e' - 'd';
        ECHO;
    }
```

After the floating point constant is recognized, it is scanned by the `for` loop to find the letter `d` or `D`. The program then adds `'e'-'d'`, which converts it to the next letter of the alphabet. The modified constant, now single-precision, is written out again. There follow a series of names which must be respelled to remove their initial `d`. By using the array `yytext` the same action suffices for all the names (only a sample of a rather long list is given here).

```
{d}{s}{i}{n}      |
{d}{c}{o}{s}      |
{d}{s}{q}{r}{t}   |
{d}{a}{t}{a}{n}   |
...
{d}{f}{l}{o}{a}{t} printf("%s",yytext+1);
```

Another list of names must have initial `d` changed to initial `a`:

```
{d}{l}{o}{g}      |
{d}{l}{o}{g}10    |
{d}{m}{i}{n}1     |
{d}{m}{a}{x}1     {
    yytext[0] += 'a' - 'd';
    ECHO;
}
```

And one routine must have initial `d` changed to initial `r`:

```

{d}1{m}{a}{c}{h}      {yytext[0] =+ 'r' - 'd';
                        ECHO;
                        }

```

To avoid such names as *dsinx* being detected as instances of *dsin*, some final rules pick up longer words as identifiers and copy some surviving characters:

```

[A-Za-z] [A-Za-z0-9]* |
[0-9]+ |
\n |
.      ECHO;

```

Note that this program is not complete; it does not deal with the spacing problems in FORTRAN or with the use of keywords as identifiers.

7.9. Left Context-Sensitivity

Sometimes it is desirable to have several sets of lexical rules to be 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, and there are several ways of handling such problems. The \wedge operator, for example, is a prior context operator, recognizing immediately preceding left context just as $\$$ recognizes immediately following right context. Adjacent left context could be extended, to produce a facility similar to that for adjacent right context, but it is unlikely to be as useful, since often the relevant left context appeared some time earlier, such as at the beginning of a line.

This section describes three means of dealing with different environments: a simple use of flags, when only a few rules change from one environment to another, the use of *start conditions* on rules, and the possibility of making multiple lexical analyzers all run together. In each case, there are rules which recognize the need to change the environment in which the following input text is analyzed, and set some parameter to reflect the change. This may be a flag explicitly tested by the programmer's action code; such a flag is the simplest way of dealing with the problem, since `lex` is not involved at all. It may be more convenient, however, to have `lex` remember the flags as initial conditions on the rules. Any rule may be associated with a start condition. It is only be recognized when `lex` is in that start condition. The current start condition may be changed at any time. Finally, if the sets of rules for the different environments are very dissimilar, clarity may be best achieved by writing several distinct lexical analyzers, and switching from one to another as desired.

Consider the following problem: copy the input to the output, changing the word *magic* to *first* on every line which begins with the letter *a*, changing *magic* to *second* on every line which begins with the letter *b*, and changing *magic* to *third* on every line which 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:

```

    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;
    }
}

```

should be adequate.

To handle the same problem with start conditions, each start condition must be introduced to `lex` in the definitions section with a line reading

```
%Start name1 name2 ...
```

where the conditions may be named in any order. The word *Start* may be abbreviated to *s* or *S*. The conditions may be referenced at the head of a rule with the `<>` brackets:

```
<name1>expression
```

is a rule which is only recognized when `lex` is in the start condition *name1*. To enter a start condition, execute the action statement

```
BEGIN name1;
```

which changes the start condition to *name1*. To resume the normal state,

```
BEGIN 0;
```

which resets to the initial condition of the `lex` automaton interpreter. A rule may be active in several start conditions:

```
<name1,name2,name3>
```

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

The same example as before can be written:

```

%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");

```

where the logic is exactly the same as in the previous method of handling the problem, but `lex` does the work rather than the programmer's code.

7.10. Character Set

The programs generated by `lex` handle character I/O only through the routines `input`, `output`, and `unput`. Thus the character representation provided in these routines is accepted by `lex` and employed to return values in `yytext`. For internal use a character is represented as a small integer which, if the standard library is used, has a value equal to the integer value of the bit pattern representing the character on the host computer. Normally, the letter `a` is represented in the same form as the character constant `'a'`. If this interpretation is changed, by providing I/O routines which translate the characters, `lex` must be told about it, by giving a translation table. This table must be in the definitions section, and must be bracketed by two lines containing only `%T`. The table contains lines of the form

```
{integer} {character string}
```

which indicate the value associated with each character. Thus the next example

```
%T
 1 Aa
 2 Bb
...
26 Zz
27 \n
28 +
29 -
30 0
31 1
...
39 9
%T
```

Figure 7-3 *Sample character table.*

maps the lower and upper case letters together into the integers 1 through 26, newline into 27, + and - into 28 and 29, and the digits into 30 through 39. Note the escape for newline. If a table is supplied, every character that is to appear either in the rules or in any valid input must be included in the table. No character may be assigned the number 0, and no character may be assigned a bigger number than the size of the hardware character set.

7.11. Summary of Source Format

The general form of a `lex` source file is:

```
{definitions}
%%
{rules}
%%
{programmer subroutines}
```

The definitions section contains a combination of

- 1) Definitions, in the form 'name space translation'.
- 2) Included code, in the form 'space code'.
- 3) Included code, in the form

```
%{
code
%}
```

- 4) Start condition declarations, given in the form

```
%S name1 name2 ...
```

- 5) Character set tables, in the form

```
%T
number space character-string
...
%T
```

- 6) 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:

Table 7-1 *Changing Internal Array Sizes in lex*

<i>Letter</i>	<i>Parameter</i>
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 may be continued on succeeding lines by using braces to delimit it.

Regular expressions in `lex` use the following operators:

Table 7-2 *Regular Expression Operators in lex*

<i>Operator</i>	<i>Meaning</i>
<code>x</code>	the character "x"
<code>"x"</code>	an "x", even if x is an operator
<code>\x</code>	an "x", even if x is an operator
<code>[xy]</code>	the character x or y
<code>[x-z]</code>	the characters x, y or z
<code>[^x]</code>	any character but x
<code>.</code>	any character but newline
<code>^x</code>	an x at the beginning of a line
<code><y>x</code>	an x when <i>lex</i> is in start condition y
<code>x\$</code>	an x at the end of a line
<code>x?</code>	an optional x
<code>x*</code>	0,1,2, ... instances of x
<code>x+</code>	1,2,3, ... instances of x
<code>x y</code>	an x or a y
<code>(x)</code>	an x
<code>x</code>	y
<code>{xx}</code>	the translation of xx from the definitions section
<code>x{m, n}</code>	m through n occurrences of x

7.12. Caveats and Bugs

There are pathological expressions which produce exponential growth of the tables when converted to deterministic automata; fortunately, they are rare.

REJECT does not rescan the input; instead it remembers the results of the previous scan. This means that if a rule with trailing context is found, and REJECT is executed, the programmer must not have used *unput* to change the characters forthcoming from the input stream. This is the only restriction on the programmer's ability to manipulate the not-yet-processed input.

Yacc — Yet Another Compiler-Compiler

Yacc — Yet Another Compiler-Compiler	143
8.1. Basic Specifications	145
8.2. Actions	147
8.3. Lexical Analysis	149
8.4. How the Parser Works	151
8.5. Ambiguity and Conflicts	155
8.6. Precedence	159
8.7. Error Handling	162
8.8. The Yacc Environment	163
8.9. Hints for Preparing Specifications	164
Input Style	164
Left Recursion	165
Lexical Tie-ins	165
Reserved Words	166
8.10. Advanced Topics	166
Simulating Error and Accept in Actions	167
Accessing Values in Enclosing Rules.	167
Support for Arbitrary Value Types	167
8.11. A Simple Example	170
8.12. Yacc Input Syntax	173
8.13. An Advanced Example	176
8.14. Old Features Supported but not Encouraged	182

Yacc — Yet Another Compiler-Compiler

Computer program input generally has some structure; in fact, every computer program that does input can be thought of as defining an 'input language' which it accepts. An input language may be as complex as a programming language, or as simple as a sequence of numbers. Unfortunately, usual input facilities are limited, difficult to use, and often are lax about checking their inputs for validity.

`yacc` provides a general tool for describing the input to a computer program. The `yacc` programmer specifies the structure of the input, together with code to be invoked as each item is recognized. `yacc` turns such a specification into a subroutine that handles the input process; frequently, it is convenient and appropriate to have most of the flow of control in the programmer's application handled by this subroutine.

The input subroutine produced by `yacc` calls a programmer-supplied routine to return the next basic input item. Thus, the programmer can specify his input in terms of individual input characters, or in terms of higher-level constructs such as names and numbers. The programmer-supplied routine may also handle idiomatic features such as comment and continuation conventions, which typically defy easy grammatical specification.

The class of specifications that `yacc` accepts is a very general one: LALR(1) grammars with disambiguating rules.

In addition to compilers for C, FORTRAN, APL, Pascal, Ratfor, etc., `yacc` has also been used for less conventional languages, including a phototypesetter language, several desk calculator languages, a document retrieval system, and a FORTRAN debugging system.

`yacc` provides a general tool for imposing structure on the input to a computer program. The `yacc` programmer prepares a specification of the input process; this 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 programmer-supplied low-level input routine (the *lexical analyzer*) to pick up the basic items (called *tokens*) from the input stream. These tokens are organized according to the input structure rules, called *grammar rules*; when one of these rules has been recognized, then programmer code supplied for this rule, an *action*, is invoked; actions have the ability to return values and make use of the values of other actions.

`yacc` generates its actions and output subroutines in C. Moreover, many of the syntactic conventions of `yacc` follow C.

The heart of the `yacc` 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 ;
```

Here, *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 — implying that the comma is to appear literally in the input. The colon and semicolon merely serve as punctuation in the rule, and have no significance in controlling the input. Thus, with proper definitions, the input

```
July 4, 1776
```

might be matched by the above rule.

An important part of the input process is carried out by the lexical analyzer. This routine reads the input stream, recognizing 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 referred to as *tokens*.

There is considerable leeway in deciding whether to recognize structures using the lexical analyzer or grammar rules. For example, the rules

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

might be used in the above example. The lexical analyzer would only need to recognize individual letters, and *month_name* would be a nonterminal symbol. Such low-level rules tend to waste time and space, and may complicate the specification beyond `yacc`'s ability to deal with it. Usually, the lexical analyzer would recognize the month names, and return an indication that a *month_name* was seen; in this case, *month_name* would be a token.

Literal characters such as `,` must also be passed through the lexical analyzer, and are also considered tokens.

Specification files are very flexible. It is relatively easy to add to the above example the rule

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

allowing

```
7 / 4 / 1776
```

as a synonym for

July 4, 1776

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 may 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 may be self-contradictory, or they may 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 which 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.

The theory underlying `yacc` has been described elsewhere[2], [3], [4].

The next several sections describe the basic process of preparing a `yacc` specification; Section 8.1 describes the preparation of grammar rules, Section 8.2 the preparation of the programmer-supplied actions associated with these rules, and Section 8.3 the preparation of lexical analyzers. Section 8.4 describes the operation of the parser. Section 8.5 discusses various reasons why `yacc` may be unable to produce a parser from a specification, and what to do about it. Section 8.6 describes a simple mechanism for handling operator precedences in arithmetic expressions. Section 8.7 discusses error detection and recovery. Section 8.8 discusses the operating environment and special features of the parsers `yacc` produces. Section 8.9 gives some suggestions which should improve the style and efficiency of the specifications. Section 8.10 discusses some advanced topics. Section 8.11 has a brief example, and section 8.12 gives a summary of the `yacc` input syntax. Section 8.13 gives an example using some of the more advanced features of `yacc`, and, finally, section 8.14 describes mechanisms and syntax no longer actively supported, but provided for historical continuity with older versions of `yacc`.

8.1. Basic Specifications

Names refer to either tokens or nonterminal symbols. `yacc` requires token names to be declared as such. In addition, for reasons discussed in Section 8.3, it is often desirable to include the lexical analyzer as part of the specification file; it may be useful to include other programs as well. Thus, every specification file consists of three sections: the *declarations*, the *(grammar) rules*, and the *programs*. The sections are separated by double percent `%%` marks. The percent `%` is generally used in `yacc` specifications as an escape character.

In other words, a full specification file looks like

```
declarations
%%
rules
%%
programs
```

The declaration section may be empty. Moreover, if the programs section is omitted, the second %% mark may be omitted also; thus, the smallest legal yacc specification is

```
%%
rules
```

Spaces (also called blanks), tabs, and newlines are ignored except that they may not appear in names or multi-character reserved symbols. Comments may appear wherever a name is legal — they are enclosed in /* . . . */, as in C and PL/I.

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

```
A : BODY ;
```

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 may be of arbitrary length, and may be made up of letters, dot '.', underscore '_', and non-initial digits. Upper and lower case letters are distinct. The names used in the body of a grammar rule may represent tokens or nonterminal symbols.

A literal consists of a character enclosed in single quotes ''. As in C, the backslash '\' is an escape character within literals, and all the C escapes are recognized. Thus

```
'\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 NUL character ('\0' or 0) should never be used in grammar rules.

If there are several grammar rules with the same left hand side, the vertical bar '|' can be used to avoid rewriting the left hand side. In addition, 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 as

```
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 more readable, and easier to change.

If a nonterminal symbol matches the empty string, this can be indicated in the obvious way:

```
empty :   ;
```

Names representing tokens must be declared; this is most simply done by writing

```
%token  name1 name2 . . .
```

in the declarations section. See Sections 3, 5, and 6 for much more discussion. Every name not defined in the declarations section is assumed to represent a non-terminal symbol. Every nonterminal symbol must appear on the left side of at least one rule.

Of all the nonterminal symbols, one, called the *start symbol*, has particular importance. The parser is designed to recognize the 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 hand side of the first grammar rule in the rules section. It is possible, and in fact desirable, to declare the start symbol explicitly in the declarations section using the `%start` keyword:

```
%start  symbol
```

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

It is the job of the programmer-supplied lexical analyzer to return the endmarker when appropriate — see Section 8.3, below. Usually the endmarker represents some reasonably obvious I/O status, such as ‘end-of-file’ or ‘end-of-record’.

8.2. Actions

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

An action is an arbitrary C 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 curly braces ‘{’ and ‘}’. For example,

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

and

```

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

```

are grammar rules with actions.

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 pseudo-variable '\$\$' to some value. For example, an action that does nothing but return the value 1 is

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

To obtain the values returned by previous actions and the lexical analyzer, the action may use the pseudo-variables \$1, \$2, . . . , which refer to the values returned by the components of the right side of a rule, reading from left to right. Thus, if the rule is

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

for example, then \$2 has the value returned by C, and \$3 the value returned by D.

As a more concrete example, consider the rule

```
expr   :   '(' expr ')' ;
```

The value returned by this rule is usually the value of the *expr* in parentheses. This can be indicated by

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

By default, the value of a rule is the value of \$1 (the first element in it). Thus, grammar rules of the form

```
A      :      B ;
```

frequently need not have an explicit action.

In the examples above, all the actions came at the end of their rules. Sometimes, it is desirable to get control before a rule is fully parsed. yacc permits an action to be written in the middle of a rule as well as at the end. This rule is assumed to return a value, accessible through the usual \$ mechanism by the actions to the right of it. In turn, it may access the values returned by the symbols to its left. Thus, in the rule

```

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

```

the effect is to set *x* to 1, and *y* to the value returned by C.

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 off by recognizing this

added rule. yacc actually treats the above example as if it had been written:

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

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

In many applications, output is not done directly by the actions; rather, 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. For example, suppose there is a C function `node`, written so that the call

```
node( L, n1, n2 )
```

creates a node with label `L`, and descendants `n1` and `n2`, and returns the index of the newly created node. The parse tree can be built by supplying actions such as:

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

in the specification.

The programmer may 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 in `'yy'`; the programmer should avoid such names.

In these examples, all the values are integers: a discussion of values of other types will be found in Section 8.10.

8.3. Lexical Analysis

The programmer 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 in order for communication between them to take place. The numbers may be chosen by yacc, or chosen by the programmer. In either case, the `'# define'` mechanism of C is used to allow the lexical analyzer to return these numbers symbolically. For example, suppose that the token name `DIGIT` has been defined in the declarations section of the yacc specification file. The relevant portion of the lexical analyzer might look like:

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

This mechanism leads to clear, easily modified lexical analyzers; the only pitfall is the need to avoid using any token names in the grammar that are reserved or significant in C or the parser; for example, the use of `if` or `while` as token names will almost certainly cause severe difficulties when the lexical analyzer is compiled. The token name `error` is reserved for error handling, and should not be used naively (see Section 8.7).

As mentioned above, the token numbers may be chosen by `yacc` or by the programmer. In the default situation, the numbers are chosen by `yacc`. The default token number for a literal character is the numerical 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 nonnegative 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 definition. It is important that all token numbers be distinct.

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

A very useful tool for constructing lexical analyzers is the *lex* program developed by Mike Lesk⁸ and described in the previous chapter of this manual. These lexical analyzers are designed to work in close harmony with `yacc` parsers. The specifications use regular expressions instead of grammar rules. *Lex* can be easily used to produce quite complicated lexical analyzers, but there remain some languages (such as FORTRAN) which do not fit any theoretical framework, and whose lexical analyzers must be crafted by hand.

8.4. How the Parser Works

`yacc` turns the specification file into a C program, which parses the input according to the specification given. The algorithm used to go from the specification to the parser is complex, and will not be discussed here (see the references for more information). The parser itself, however, is relatively simple, and understanding how it works, while not strictly necessary, will nevertheless make 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 can read and remember the next input token (called the *lookahead* 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 lookahead token has been read.

The machine has only four actions available to it, called *shift*, *reduce*, *accept*, and *error*. A move of the parser is done as follows:

1. Based on its current state, the parser decides whether it needs a lookahead token to decide what action should be done; if it needs one, and does not have one, it calls `yylex` to obtain the next token.
2. Using the current state, and the lookahead token if needed, the parser decides on its next action, and carries it out. This may result in states being pushed onto the stack, or popped off the stack, and in the lookahead token being processed or left alone.

shift Action

The *shift* action is the most common action the parser takes. Whenever a shift action is taken, there is always a lookahead token. For example, in state 56 there may be an action:

```
IF      shift 34
```

which says, in state 56, if the lookahead 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). The lookahead token is cleared.

reduce Action

The *reduce* action keeps the stack from growing without bound. Reduce actions are appropriate when the parser has seen the right hand side of a grammar rule, and is prepared to announce that it has seen an instance of the rule, replacing the right hand side by the left hand side. It may be necessary to consult the lookahead token to decide whether to reduce, but usually it is not; in fact, the default action (represented by a `.`) is often a reduce action.

Reduce actions are associated with individual grammar rules. Grammar rules are also given small integer numbers, leading to some confusion. The action

```
. reduce 18
```

refers to *grammar rule* 18, while the action

```
IF      shift 34
```

refers to *state* 34.

Suppose the rule being reduced is

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

The reduce action depends on the left hand symbol (A in this case), and the number of symbols on the right hand 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 which 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, pushed onto the stack, and parsing continues. There are significant differences between the processing of the left hand symbol and an ordinary shift of a token, however, so this action is called a *goto* action. In particular, the lookahead token is cleared by a shift, and is not affected by a *goto*. In any case, the uncovered state contains an entry such as:

```
A      goto 20
```

which pushes state 20 onto the stack, and becomes the current state.

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 hand 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 hand side of the rule is empty, no states are popped off the stack: the uncovered state is in fact the current state.

The reduce action is also important in the treatment of programmer-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 `yyval` is copied onto the value stack. After the return from the programmer’s code, the reduction is carried out. When the *goto* action is done, the external variable `yyval` is copied onto the value stack. The pseudo-variables \$1, \$2, etc., refer to the value stack.

accept and error Actions

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 lookahead token is the endmarker, and indicates that the parser has successfully done 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 lookahead 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) will be covered in Section 8.7.

It is time for an example! Consider the 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 `y.output` file corresponding to the above grammar (with some statistics stripped off the end) is:

```
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
```

Notice that, in addition to the actions for each state, there is a description of the parsing rules being processed in each state. The `_` character is used to indicate what has been seen, and what is yet to come, in each rule. Suppose the input is

```
DING DONG DELL
```

It is instructive to follow the steps of the parser while processing this input.

Initially, the current state is state 0. The parser needs to refer to the input in order to decide between the actions available in state 0, so the first token, DING,

is read, becoming the lookahead token. The action in state 0 on DING is ‘shift 3’, so state 3 is pushed onto the stack, and the lookahead token is cleared. State 3 becomes the current state. The next token, DONG, is read, becoming the lookahead token. The action in state 3 on the token DONG is ‘shift 6’, so state 6 is pushed onto the stack, and the lookahead is cleared. The stack now contains 0, 3, and 6. In state 6, without even consulting the lookahead, the parser reduces by rule 2.

```
sound : DING DONG
```

This rule has two symbols on the right hand side, so 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*,

```
sound goto 2
```

is obtained; thus state 2 is pushed onto the stack, becoming 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 lookahead token is cleared. In state 5, the only action is to reduce by rule 3. This has one symbol on the right hand 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; the endmarker is obtained, indicated by ‘\$end’ in the *y.output* file. The action in state 1 when the endmarker is seen is to accept, successfully ending 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 spend with this and other simple examples will probably be repaid when problems arise in more complicated contexts.

8.5. 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 grammar rule

```
expr : expr '-' expr
```

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. Unfortunately, this grammar rule does not unambiguously 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 as

```
expr - ( expr - expr )
```

The first is called *left association*, the second *right association*.

`yacc` detects such ambiguities when it is attempting to build the parser. It is instructive to consider the problem that confronts the parser when it is given an input such as

```
expr - expr - expr
```

When the parser has read the second `expr`, the input that it has seen:

```
expr - expr
```

matches the right side of the grammar rule above. The parser could *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 would then read the final part of the input:

```
- expr
```

and again reduce. The effect of this is to take the left-associative interpretation.

Alternatively, when the parser has seen

```
expr - expr
```

it could defer the immediate application of the rule, and continue reading the input until it had seen

```
expr - expr - expr
```

It could then apply the rule to the rightmost three symbols, reducing them to *expr* and leaving

```
expr - expr
```

Now the rule can be reduced once more; the effect is to take the right associative interpretation. Thus, having read

```
expr - expr
```

the parser can do two legal things, a shift or a reduction, and has no way of deciding between them. This is called a *shift / reduce conflict*. It may 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 which choice to make in a given situation is called a *disambiguating rule*.

`yacc` invokes two disambiguating rules by default:

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

Rule 1 implies that reductions are deferred whenever there is a choice, in favor of shifts. Rule 2 gives the programmer rather crude control over the behavior of the parser in this situation, but reduce/reduce conflicts should be avoided whenever possible.

Conflicts may 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. Our experience has suggested that this rewriting is somewhat unnatural, and produces slower parsers; thus, yacc will produce parsers even in the presence of conflicts.

As an example of the power of disambiguating rules, consider a fragment from a programming language involving an ‘if-then-else’ construction:

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

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 will be called the *simple-if rule*, and the second the *if-else rule*.

These two rules form an ambiguous construction, since input of the form

```
IF ( condition-1 ) IF ( condition-2 ) statement-1 ELSE statement-2
```

can be structured according to these rules in two ways:

```
IF ( condition-1 ) {
    IF ( condition-2 ) statement-1
}
ELSE statement-2
```

or

```
IF ( condition-1 ) {
    IF ( condition-2 ) statement-1
    ELSE statement-2
}
```

The second interpretation is the one given in most programming languages having this construct. Each ELSE is associated with the last preceding ‘un-ELSE’ d’ IF. In this example, consider the situation where the parser has seen

```
IF ( condition-1 ) IF ( condition-2 ) statement-1
```

and is looking at the ELSE. It can immediately reduce by the simple-if rule to get

```
IF ( condition-1 ) stat
```

and then read the remaining input,

ELSE *statement-2*

and reduce

IF (*condition-1*) stat ELSE *statement-2*

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

On the other hand, the ELSE may be shifted, *statement-2* read, and then the right hand portion of

IF (*condition-1*) IF (*condition-2*) *statement-1* ELSE *statement-2*

can be reduced by the if-else rule to get

IF (*condition-1*) stat

which can be reduced by the simple-if rule. 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 already seen, such as

IF (*condition-1*) IF (*condition-2*) *statement-1*

In general, there may be many conflicts, and each one will be 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
```

The first line describes the conflict, giving the state and the input symbol. The ordinary state description follows, giving the grammar rules active in the state, and the parser actions. Recall that the underline marks the portion of the grammar rules which 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 possible things. If the input symbol is ELSE, it is possible to shift into state 45.

State 45 will have, as part of its description, the line

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

since the ELSE will have been shifted in this state. Back in state 23, the alternative action, described by '.', is to be done if the input symbol is not mentioned explicitly in the above actions; thus, in this case, if the input symbol is not ELSE, the parser reduces by grammar rule 18:

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

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 which can be reduced. In most states, there will be at most one reduce action possible in the state, and this will be the default command. Programmers who encounter unexpected shift/reduce conflicts will probably want to look at the verbose output to decide whether the default actions are appropriate. In really tough cases, the programmer might need to know more about the behavior and construction of the parser than can be covered here. In this case, one of the theoretical references [2], [3], [4] might be consulted; the services of a local guru might also be appropriate.

8.6. 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 form

```
expr : expr OP expr
```

and

```
expr : UNARY expr
```

for all binary and unary operators desired. This creates a very ambiguous grammar, with many parsing conflicts. As disambiguating rules, the programmer 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 a yacc keyword: %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. Thus,

```
%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 `.LT.` operator in FORTRAN, that may not associate with themselves; thus,

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

is illegal in FORTRAN, and such an operator would be described with the keyword `%nonassoc` in `yacc`. As an example of the behavior of these declarations, the description

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

%%

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

might be used to structure the input

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

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 '-'; unary minus may 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. It changes the precedence of the grammar rule to become that of the following token name or literal. For example, to make unary minus have the same precedence as multiplication the rules might resemble:

```

%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 may 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:

1. The precedences and associativities are recorded for those tokens and literals that have them.
2. 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 may have no precedence and associativity associated with them.
3. 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.
4. 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 may disguise errors in the input grammar; it is a good idea to be sparing with precedences, and use them in an essentially ‘cook-book’ 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.

8.7. 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 may 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 programmer 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 lookahead token, and performs the action encountered. The lookahead token is then reset to the token that caused the error. If no special error rules have been specified, the processing halts when an error is detected.

In order to prevent a cascade of error messages, the parser, after detecting an error, remains in error state until three tokens have been 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 form

```
stat      : error
```

would, in effect, mean that on a syntax error the parser would attempt to skip over the statement in which the error was seen. More precisely, the parser will scan ahead, looking for three tokens that might legally follow a statement, and start processing at the first of these; if the beginnings of statements are not sufficiently distinctive, it may 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 may be used with these special error rules. These actions might attempt to reinitialize tables, reclaim symbol table space, etc.

Error rules such as the above are very general, but difficult to control. Somewhat easier are rules such as

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

Here, when there is an error, the parser attempts to skip over the statement, but will do so by skipping to the next ';'. All tokens after the error and before the next ';' cannot be shifted, and are discarded. When the ';' is seen, this rule will be reduced, and any 'cleanup' action associated with it performed.

Another form of error rule arises in interactive applications, where it may be desirable to permit a line to be reentered after an error. A possible error rule might be

```
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 has 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 be used to force the parser to believe that an error has been fully recovered from. The statement

```
yyerrok ;
```

in an action resets the parser to its normal mode. The last example is better written

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

As mentioned above, 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 lookahead token must be cleared. The statement

```
yyclearin ;
```

in an action will have this effect. For example, suppose the action after error were to call some sophisticated resynchronization routine, supplied by the programmer, that attempted to advance the input to the beginning of the next valid statement. After this routine was called, the next token returned by `yylex` would presumably be the first token in a legal statement; the old, illegal token must be discarded, and the error state reset. This could be done by a rule like

```
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; moreover, the programmer can get control to deal with the error actions required by other portions of the program.

8.8. The Yacc Environment

When the programmer inputs a specification to `yacc`, the output is a file of C programs, called `y.tab.c` on most systems (due to local file system conventions, the name may differ from installation to installation). `yacc` produces an integer-valued function called `yyparse`. When `yyparse` is called, it in turn repeatedly calls `yylex` — the lexical analyzer supplied by the programmer (see Section 8.3) to obtain input tokens. Eventually, either an error is detected, in which case (if no error recovery is possible) `yyparse` returns the value 1, or the

lexical analyzer returns the endmarker token and the parser accepts. In this case, `yyparse` returns the value 0.

The programmer must provide a certain amount of environment for this parser in order to obtain a working program. For example, as with every C program, a program called `main` must be defined, that eventually calls `yyparse`. In addition, a routine called `yyerror` prints a message when a syntax error is detected.

The programmer must supply these two routines in one form or another. They can be as simple as the following example, or they can be as complex as needed.

```
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 will want 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 lookahead token number at the time the error was detected; this may be of some interest in giving better diagnostics.

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

8.9. Hints for Preparing Specifications

This section contains miscellaneous hints on preparing efficient, easy to change, and clear specifications. The individual subsections are more or less independent.

Input Style

It is difficult to provide rules with substantial actions and still have a readable specification file. The following style hints owe much to Brian Kernighan.

- a. Use all capital letters for token names, all lower case letters for nonterminal names. This rule comes under the heading of 'knowing who to blame when things go wrong.'
- b. Put grammar rules and actions on separate lines. This allows either to be changed without an automatic need to change the other.
- c. Put all rules with the same left hand side together. Put the left hand side in only once, and let all following rules begin with a vertical bar.
- d. Put a semicolon only after the last rule with a given left hand side, and put the semicolon on a separate line. This allows new rules to be added easily.

- e. Indent rule bodies by two tab stops, and action bodies by three tab stops.

The example in section 8.11 is written following this style, as are the examples in the text of this paper (where space permits). The programmer must make up his 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 form

```
name      :  name  rest_of_rule  ;
```

These rules frequently arise when writing specifications of sequences and lists:

```
list      :  item
          |  list  ‘,’  item
          ;
```

and

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

In each of these cases, the first rule will be reduced for the first item only, and the second rule will be reduced for the second and all succeeding items.

With right-recursive rules, such as

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

the parser would be a bit bigger, and the items would be seen, and reduced, from right to left. More seriously, an internal stack in the parser would be in danger of overflowing if a very long sequence were read. Thus, the programmer should use left recursion wherever reasonable.

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

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

Once again, the first rule would always be reduced exactly once, before the first item was read, and then the second rule would be 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 Tie-ins

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. For example, suppose a program consists

of 0 or more declarations, followed by 0 or more statements. Consider:

```
%{
    int dflag;
%}
... other declarations ...

%%

prog  :  decls  stats
      ;

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

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

... other rules ...
```

The flag *dflag* is now 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 has ended and the statements have begun. In many cases, this single-token exception does not affect the lexical scan.

This kind of ‘backdoor’ approach can be elaborated to a noxious 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 the programmer to use words like ‘if’, which 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 programmer can make a stab at it, using the mechanism described in the last subsection, but it is difficult.

A number of ways of making this easier are under advisement. Until then, it is better that the keywords be *reserved*; that is, be forbidden for use as variable names. There are powerful stylistic reasons for preferring this, anyway.

8.10. Advanced Topics

This section discusses a number of advanced features of *yacc*.

Simulating Error and Accept in Actions

The parsing actions of error and accept can be simulated in an action by use of macros `YYACCEPT` and `YYERROR`. `YYACCEPT` makes `yyparse` return the value 0; `YYERROR` makes the parser behave as if the current input symbol results in a syntax error; `yyperror` is called, and error recovery takes place. These mechanisms can be used to simulate parsers with multiple endmarkers or context-sensitive syntax checking.

Accessing Values in Enclosing Rules.

An action may refer to values returned by actions to the left of the current rule. The mechanism is simply the same as with ordinary actions, a dollar sign followed by a digit, but in this case the digit may be 0 or negative. Consider

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

Support for Arbitrary Value Types

By default, the values returned by actions and the lexical analyzer are integers. `yacc` can also support values of other types, including structures. In addition, `yacc` keeps track of the types, and inserts appropriate union member names so that the resulting parser will be strictly type checked. The `yacc` value stack (see Section 8.4) is declared to be a union of the various types of values desired. The programmer declares the union, and associates a union member name 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 will take place. In addition, type-checking commands such as `lint(1)` will be far more silent.

There are three mechanisms used to provide for this typing. First, there is a way of defining the union; this must be done by the programmer since 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 non-terminals. Finally, there is a mechanism for describing the type of those few values where `yacc` cannot easily determine the type.

To declare the union, the programmer includes 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 may be declared in a header file, and a `typedef` used to define the variable `YYSTYPE` to represent this union. Thus, the header file might also have said:

```
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 construction

```
< name >
```

is used to indicate a union member name. If this follows one of the keywords `%token`, `%left`, `%right`, and `%nonassoc`, the union member name is associated with the tokens listed. Thus, saying

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

will tag any reference to values returned by these two tokens with the union member name `optype`. Another keyword, `%type`, is used similarly to associate union member names with nonterminals. Thus, one might say

```
%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 `$0` — see the previous subsection) leaves `yacc` with no easy way of knowing the type. In this case, a type can be imposed on the reference by inserting a union member name, between `<` and `>`, immediately after the first `$`. An example of this usage is

```
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 8.13. The facilities in this subsection are not triggered until they are used: in particular, the use of `%type` will turn 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's, as was true historically. This paper is reprinted in this manual.

8.11. A Simple Example

This example gives the complete `yacc` specification for a small desk calculator; the desk calculator has 26 registers, labeled 'a' through 'z', and accepts arithmetic expressions made up of the operators `+`, `-`, `*`, `/`, `%` (mod operator), `&` (bitwise and), `|` (bitwise or), and assignment. If an expression at the top level is an assignment, the value is not printed; otherwise it is. As in C, 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 precedences and ambiguities are used, and demonstrating simple error recovery. The major oversimplifications are that the lexical analysis phase is much simpler than for most applications, and the output is produced immediately, line-by-line. Note the way that decimal and octal integers are read in by the 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 /* supplies precedence for unary minus */

%% /* beginning of rules section */

list : /* empty */
     | list stat '\n'
     | list error '\n'
       { yyerrok; }
     ;

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
       { $$ = $1 / $3; }
     | expr '%' expr
       { $$ = $1 % $3; }
     | expr '&' expr
       { $$ = $1 & $3; }
     | expr '|' expr
       { $$ = $1 | $3; }
     | '-' expr %prec UMINUS
       { $$ = - $2; }
     | LETTER

```

```
        { $$ = regs[$1]; }
    | number
    ;

number : DIGIT
       { $$ = $1;   base = ($1==0) ? 8 : 10; }
    | number DIGIT
       { $$ = base * $1 + $2; }
    ;

%% /* start of programs */

yylex() /* lexical analysis routine */
{
/* returns LETTER for lower case letter, yylval=0 thru 25 */
/* return DIGIT for digit, yylval=0 thru 9 */
/* all other characters are returned immediately */

    int c;

    while((c = getchar()) == ' ') { /* skip blanks */ }

        /* c is now nonblank */
    if(islower(c)) {
        yylval = c - 'a';
        return(LETTER);
    }
    if(isdigit(c)) {
        yylval = c - '0';
        return(DIGIT);
    }
    return(c);
}
```


8.12. Yacc Input Syntax

This section describes the yacc input syntax, as a yacc specification. Context dependencies, etc., 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 decide whether the next token (skipping blanks, newlines, comments, etc.) is a colon. If so, it returns the token `C_IDENTIFIER`. Otherwise, it returns `IDENTIFIER`. Literals (quoted strings) are also returned as `IDENTIFIERS`, but never as part of `C_IDENTIFIER`s.

```

        /* grammar for the input to Yacc */

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

        /* reserved words: %type => TYPE, %left => LEFT, etc. */

%token LEFT RIGHT NONASSOC TOKEN PREC TYPE START UNIO

%token MARK /* the %% mark */
%token LCURL /* the %{ mark */
%token RCURL /* the %} mark */

        /* ascii character literals stand for themselves */

%start spec

%%

spec : defs MARK rules tail
      ;

tail : MARK { In this action, eat up the rest of the file }
      | /* empty: the second MARK is optional */
      ;

defs : /* empty */
      | defs def
      ;

def : 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

```

```

;

nmno  :  IDENTIFIER /* NOTE: literal illegal with %type */
      |  IDENTIFIER NUMBER /* NOTE: illegal with %type */
;

/* rules section */

rules :  C_IDENTIFIER rbody prec
      |  rules rule
;

rule  :  C_IDENTIFIER rbody prec
      |  '|' rbody prec
;

rbody :  /* empty */
      |  rbody IDENTIFIER
      |  rbody act
;

act   :  '{' { Copy action, translate $$, etc. } "'"
;

prec  :  /* empty */
      |  PREC IDENTIFIER
      |  PREC IDENTIFIER act
      |  prec ';'
;
```

8.13. An Advanced Example

This section gives an example of a grammar using some of the advanced features discussed in Section 8.10. The desk calculator example in section 8.11 is modified to provide a desk calculator that does floating point interval arithmetic. The calculator understands floating point constants, the arithmetic operations $+$, $-$, $*$, $/$, unary $-$, and $=$ (assignment), and has 26 floating point variables, 'a' through 'z'. Moreover, it also understands *intervals*, written

$$(x , y)$$

where x is less than or equal to y . There are 26 interval-valued variables 'A' through 'Z' that may also be used. The usage is similar to that in section 8.11 — assignments return no value, and print nothing, while expressions print the (floating or interval) value.

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

It is also worth noting the use of `YYERROR` to handle error conditions: division by an interval containing 0, and an interval presented in the wrong order. In effect, 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. Note that a scalar 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 two 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-valued expression in the second example, but this fact is not known until the ',' 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 will be applied automatically. Despite this evasion, there are still many cases where the conversion may 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 conflicts will be 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 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 ];  
  
%}  
  
%start    lines  
  
%union    {  
    int ival;  
    double dval;  
    INTERVAL vval;  
    }  
  
%token <ival> DREG VREG /* indices into dreg, vreg arrays */  
  
%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.lo, $1.hi ); }  
    |    DREG '=' dexp '\n'
```

```

        { dreg[$1] = $3; }
| VREG '=' vexp '\n'
        { vreg[$1] = $3; }
| error '\n'
        { yyerrok; }
;

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; }
| '(' dexp ')'
        { $$ = $2; }
;

vexp : 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 = $1hi + $3hi;
          $$lo = $1lo + $3lo; }
| dexp '+' vexp
        {
          $$hi = $1 + $3hi;
          $$lo = $1 + $3lo; }
| vexp '-' vexp
        {
          $$hi = $1hi - $3lo;
          $$lo = $1lo - $3hi; }
| dexp '-' vexp
        {
          $$hi = $1 - $3lo;
          $$lo = $1 - $3hi; }
| vexp '*' vexp
        { $$ = vmul( $1lo, $1hi, $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, $1, $3 ); }
| '-' vexp %prec UMINUS
|   { $$ .hi = -$2.lo;   $$ .lo = -$2.hi;   }
| '(' vexp ')'
|   { $$ = $2; }
;

%%

# define BSZ 50 /* buffer size for floating point numbers */

/* lexical analysis */

yylex(){
  register c;

  while( (c=getchar()) == ' ' ){ /* skip over blanks */ }

  if( isupper( c ) ){
    yylval.ival = c - 'A';
    return( VREG );
  }
  if( islower( c ) ){
    yylval.ival = c - 'a';
    return( DREG );
  }

  if( isdigit( c ) || c=='.' ){
    /* gobble up digits, points, exponents */

    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 ) return( '.' ); /* will cause syntax error */
        continue;
      }

      if( c == 'e' ){
        if( exp++ ) return( 'e' ); /* will cause syntax error */
        continue;
      }

      /* end of number */

```



```

        break;
    }
    *cp = '\0';
    if( (cp-buf) >= BSZ ) printf( "constant too long: truncated\n" );
    else ungetc( c, stdin ); /* push back last char read */
    yylval.dval = atof( buf );
    return( CONST );
}
return( c );
}

INTERVAL hilo( a, b, c, d ) double a, b, c, d; {
/* returns the smallest interval containing a, b, c, and d */
/* used by *, / routines */
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 interval 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 ) );
}

```

8.14. Old Features Supported but not Encouraged

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

1. Literals may also be delimited by double quotes `""`.
2. Literals may 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 literals.

The use of multi-character literals is likely to mislead those unfamiliar with `yacc`, since it suggests that `yacc` is doing a job which must be actually done by the lexical analyzer.

3. Most places where `%` is legal, backslash `\` may be used. In particular, `\\` is the same as `%%`, `\left` the same as `%left`, etc.
4. There are a number of other synonyms:

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

5. Actions may also have the form

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

and the curly braces can be dropped if the action is a single C statement.

6. C code between `%{` and `%}` used to be permitted at the head of the rules section, as well as in the declaration section.

A

SCCS Low-Level Commands

SCCS Low-Level Commands	185
A.1. Low Level SCCS For Beginners	185
Terminology	185
A.2. SCCS File Numbering Conventions	186
A.3. Summary of SCCS Commands	188
A.4. SCCS Command Conventions	189
Flags	190
Real/Effective User	190
Back-up Files Created During Processing	190
Diagnostics	190
A.5. <code>admin</code> — Create and Administer SCCS Files	190
<code>admin</code> Options	191
Flags In SCCS Files	193
Examples of Using <code>admin</code>	195
A.6. <code>cdc</code> — Change Delta Commentary	197
<code>cdc</code> Options	197
A.7. <code>comb</code> — Combine SCCS Deltas	198
<code>comb</code> Options	199
A.8. <code>delta</code> — Make a Delta	200
<code>delta</code> Options	200
Examples of Using <code>delta</code>	202
More Notes on <code>delta</code>	202
A.9. <code>get</code> — Get Version of SCCS File	204

get Options	205
Identification Keywords	208
A.10. help — Ask for SCCS Help	216
A.11. prs — Print SCCS File	217
prs Options	218
Data Keywords	218
A.12. rmdel — Remove Delta from SCCS File	221
A.13. sact — Display SCCS Editing Activity	223
A.14. sccsdiff — Display Differences in SCCS Versions	223
sccsdiff Options	223
A.15. unget — Undo a Previous SCCS get	224
unget Options	224
A.16. val — Validate SCCS File	224
val Options	224
what — Identify SCCS Files	225
A.17. SCCS Files	226
Protection	226
Layout of an SCCS File	227
Auditing	228

SCCS Low-Level Commands

This appendix contains a summary of the individual SCCS commands. The user-level interface to SCCS is described in chapter 4 of this manual. In the unlikely event that you need to use the ‘raw’ commands of SCCS, here they are. Be aware that the commands described here do not make any assumptions about where the *s.file* are — you must spell it all out in excruciating detail. The individual SCCS tools are not easy to use, but they do provide extremely close control over the SCCS database files. Of particular interest are the numbering of branches, the *l-file*, which gives a description of what deltas were used on a `get`, and certain other SCCS commands.

The following topics are covered here:

- The scheme used to identify versions of text kept in an SCCS file.
- Basic information needed for day-to-day use of SCCS commands, including a discussion of the more useful arguments.
- Protection and auditing of SCCS files, including the differences between the use of SCCS by *individual* users on one hand, and *groups* of users on the other.

A.1. Low Level SCCS For Beginners

In this section, we present some basic concepts of SCCS. Examples are fragments of terminal sessions, with what you type shown in **bold typewriter font like this**, and what the terminal displays shown in typewriter font like this.

Note that all the SCCS commands described here live in the `/usr/sccs` directory, so you must either state that directory explicitly when using SCCS commands, or include that pathname in your `.login` file. All examples shown here assume that you have `/usr/sccs` in your path and so you just have to type the required SCCS command name.

Terminology

Each SCCS file is composed of one or more sets of changes applied to the null (empty) version of the file; each set of changes usually depends on all previous sets. Each set of changes is called a ‘delta’ and is assigned a name called the SCCS *ID*entification string (SID).

The SID is composed of at most four components; for now let’s focus on only the first two: the ‘release’ and ‘level’ numbers. Each set of changes to a file is named ‘*release.level*’; hence, the first delta is called ‘1.1’, the second ‘1.2’, the

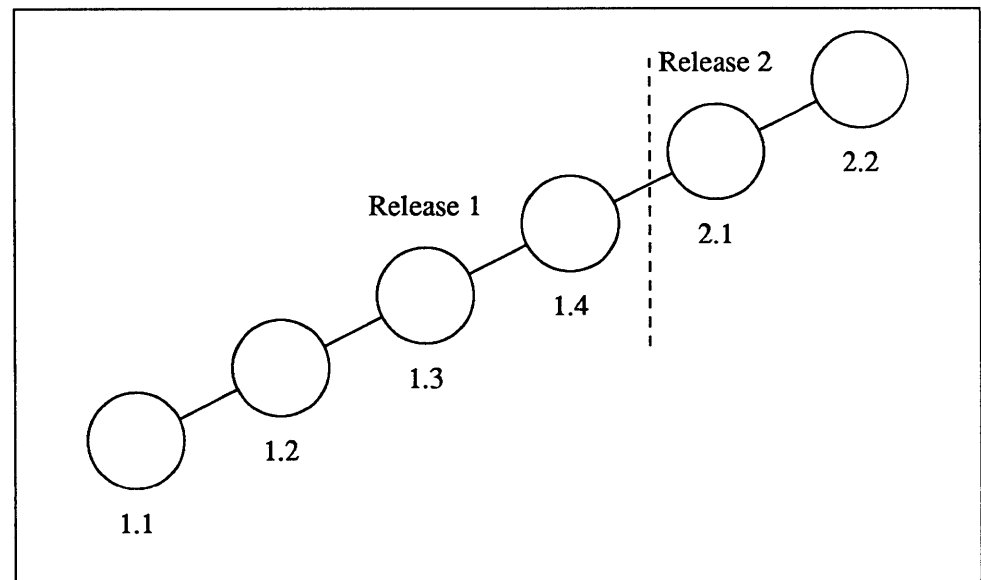
third '1.3', and so on. The release number can also be changed, allowing, for example, deltas '2.1', '3.19', etc. A change in the release number usually indicates a major change to the file.

Each delta of an SCCS file defines a particular version of the file. For example, delta 1.5 defines the version of the SCCS file obtained by applying the changes that constitute deltas 1.1, 1.2, etc., up to and including delta 1.5 itself, in that order, to the null (empty) version of the file. A.16.2.

A.2. SCCS File Numbering Conventions

You can think of the deltas applied to an SCCS file as the nodes of a tree; the root is the initial version of the file. The root delta (node) is normally named '1.1' and successor deltas (nodes) are named '1.2', '1.3', etc. We have already discussed these two components of the names of the deltas, the 'release' and 'level' numbers; and you have seen that normal naming of successor deltas proceeds by incrementing the level number, which is performed automatically by SCCS whenever a delta is made. In addition, you have seen how to change the release number when making a delta, to indicate that a major change to the file is being made. The new release number applies to all successor deltas, unless it is specifically changed again. Thus, the evolution of a particular file may be represented as in Figure A-1.

Figure A-1 *Evolution of an SCCS File*



We can call this structure the 'trunk' of the SCCS tree. It represents the normal sequential development of an SCCS file, in which changes that are part of any given delta are dependent upon all the preceding deltas.


Branches

However, there are situations when a *branch* is needed on the tree: when changes applied as part of a given delta are *not* dependent upon all previous deltas. As an example, consider a program which is in production use at version 1.3, and for which development work on release 2 is already in progress. Thus, release 2 may

already have some deltas, precisely as shown in Figure 1. Assume that a production user reports a problem in version 1.3 which cannot wait until release 2 to be repaired. The changes necessary to repair the trouble will be applied as a delta to version 1.3 (the version in production use). This creates a new version that will then be released to the user, but will *not* affect the changes being applied for release 2 (that is, deltas 1.4, 2.1, 2.2, etc.).

The new delta is a node on a 'branch' of the tree, and its name consists of four components: the release and level numbers, as with trunk deltas, plus the 'branch' and 'sequence' numbers. Its SID thus appears as: *release.level.branch.sequence*. The *branch* number is assigned to each branch that is a descendant of a particular trunk delta; the first such branch is 1, the next one 2, and so on. The *sequence* number is assigned, in order, to each delta on a particular branch. Thus, 1.3.1.2 identifies the second delta of the first branch that derives from delta 1.3. This is shown in Figure A-2.

Figure A-2 *Tree Structure with Branch Deltas*




The concept of branching may be extended to any delta in the tree; the naming of the resulting deltas proceeds in the manner just illustrated.

Two observations are of importance with regard to naming deltas. First, the names of trunk deltas contain exactly two components, and the names of branch deltas contain exactly four components. Second, the first two components of the name of a branch delta are always those of the ancestral trunk delta, and the branch component is assigned in the order of creation of the branch, independently of its location relative to the trunk delta. Thus, a branch delta may always be identified as such from its name. Although the ancestral trunk delta may be identified from the branch delta's name, it is *not* possible to determine the *entire* path leading from the trunk delta to the branch delta. For example, if delta 1.3 has one branch emanating from it, all deltas on that branch will be named 1.3.1.*n*.

If a delta on this branch then has another branch emanating from it, all deltas on the new branch will be named 1.3.2.*n* (see Figure A-3. The only information that may be derived from the name of delta 1.3.2.2 is that it is the *chronologically* second delta on the *chronologically* second branch whose *trunk* ancestor is delta 1.3. In particular, it is *not* possible to determine from the name of delta 1.3.2.2 all of the deltas between it and its trunk ancestor (1.3).

Figure A-3 *Extending the Branching Concept*



It is obvious that the concept of branch deltas allows the generation of arbitrarily complex tree structures. Although this capability has been provided for certain specialized uses, it is strongly recommended that the SCCS tree be kept as simple as possible, because comprehension of its structure becomes extremely difficult as the tree becomes more complex.

A.3. Summary of SCCS Commands

Here is a summary of all the SCCS commands and their major functions:

admin	Creates SCCS files and applies changes to parameters of SCCS files. admin is described in section A.5.
cdc	Changes the commentary associated with a delta. cdc is described in section A.6.
comb	Combines two or more consecutive deltas of an SCCS file into a single delta. comb is described in section A.7.
delta	Applies changes (deltas) to the text of SCCS files; that is, delta creates new versions. delta is described in section A.8.
get	Retrieves versions of SCCS files. get is described in section A.9.
help	Explains SCCS commands and diagnostic messages. help is described in section A.10.

<code>prs</code>	Prints portions of an SCCS file in user-specified format. <code>prs</code> is described in section A.11.
<code>rmDEL</code>	Removes a delta from an SCCS file; useful for removing deltas that were created by mistake. <code>rmDEL</code> is described in section A.12.
<code>sccsdiff</code>	Shows the differences between any two versions of an SCCS file. <code>sccsdiff</code> is described in section A.14.
<code>val</code>	Validates an SCCS file. <code>val</code> is described in section A.16.
<code>what</code>	Searches UNIX [†] file(s) for all occurrences of a special pattern and prints what follows it. <code>what</code> is useful in finding identifying information inserted by <code>get</code> . <code>what</code> is described in section

A.4. SCCS Command Conventions

This section discusses the conventions and rules that apply to SCCS commands. These rules and conventions are generally applicable to *all* SCCS commands, except as indicated below.

SCCS commands, like most UNIX commands, accept *options* and *file arguments*.

Options

Options begin with a minus sign (-), followed by a lower-case alphabetic character, and, in some cases, followed by a value. Options modify actions of commands on which they are specified.

File arguments

File arguments (which may be names of files and/or directories) specify the file(s) that the given SCCS command is to process; naming a directory is equivalent to naming *all* the SCCS files within the directory. Non-SCCS files and unreadable files in the named directories are silently ignored.

In general, file arguments may *not* begin with a minus sign. However, if the name '-' (a lone minus sign) is specified as an argument to a command, the command reads the standard input for lines and takes each line as the *name* of an SCCS file to be processed. The standard input is read until end-of-file. This feature is often used in pipelines with, for example, the `find(1)` or `ls(1)` commands. Again, names of non-SCCS files and of unreadable files are silently ignored.

Options specified for a given command apply to *all* file arguments of that command. Options are processed before any file arguments; therefore the placement of options is arbitrary, that is, options may be interspersed with file arguments. File arguments, however, are processed left to right.

Somewhat different argument conventions apply to the `help`, `what`, `sccsdiff`, and `val` commands.

[†] UNIX is a trademark of AT&T Bell Laboratories.

Flags Certain actions of various SCCS commands are modified by *flags* embedded in the text of SCCS files. Some of these flags are discussed below. For a complete description of all such flags, see `admin`.

Real/Effective User The distinction between the *real user* (see `passwd(1)`) and the *effective user* of a UNIX system is of concern in discussing various actions of SCCS commands. For the present, it is assumed that both the real user and the effective user are one and the same, that is, the user who is logged into the system.

Back-up Files Created During Processing 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 will not be damaged if processing terminates abnormally. The name of the *x-file* is formed by replacing the 's.' of the SCCS file name with 'x.'. When processing is complete, the old SCCS file is removed and the *x-file* is renamed to be the SCCS file. The *x-file* is created in the directory containing the SCCS file, is given the same mode (see `chmod(1)`) as the SCCS file, and is owned by the effective user.

To prevent simultaneous updates to an SCCS file, commands that modify SCCS files create a *lock-file*, called the *z-file*, whose name is formed by replacing the 's.' of the SCCS file name with 'z.'. The *z-file* contains the *process number* of the command that creates it, and its existence is an indication to other commands that that SCCS file is being updated. Thus, other commands that modify SCCS files will not process an SCCS file if the corresponding *z-file* exists. The *z-file* is created with mode 444 (read-only) in the directory containing the SCCS file, and is owned by the effective user. The *z-file* exists only for the duration of the execution of the command that creates it. In general, users can ignore *x-files* and *z-files*; they may be useful in the event of system crashes or similar situations.

Diagnostics SCCS commands direct their diagnostic responses to the standard error file. SCCS diagnostics generally look like this:

```
ERROR [name-of-file-being-processed]: message text (code)
```

The *code* in parentheses may be used as an argument to `help` to obtain a further explanation of the diagnostic message.

If the SCCS command detects a fatal error during the processing of a file it terminates processing of *that* file and proceeds with the next file in the series, if more than one file has been named.

A.5. `admin` — Create and Administer SCCS Files

`admin` creates new SCCS files and changes parameters of existing ones. Options and SCCS file names may appear in any order on the `admin` command line. SCCS file names must begin with the characters 's.'. A named file is created if it doesn't exist already, and its parameters are initialized according to the specified options. Any parameter not initialized by an option is assigned a default value. If a named file does exist, parameters corresponding to specified options are changed, and other parameters are left as is.

```
admin [ -n ] [ -i [ name ] ] [ -rrel ] [ -t [ name ] ] [ -f flag [ flag-val ] ] ...
[ -d flag [ flag-val ] ] ... [ -a login ] ... [ -e login ] ... [ -m [ mrlist ] ]
[ -y [ comment ] ] [ -h ] [ -z ] filename ...
```

If a directory is named, `admin` behaves as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the path name does not begin with `s.`) and unreadable files are silently ignored. A name of `-` means the standard input — each line of the standard input is taken as the name of an SCCS file to be processed. Again, non-SCCS files and unreadable files are silently ignored.

admin Options

Options are explained as though only one named file is to be processed, since options apply independently to each named file.

Creating a new file

`-n` A new SCCS file is being created.

Initial text

`-i [name]`

Initial text: file *name* contains the text of a new SCCS file. The text is the first delta of the file — see `-r` option for delta numbering scheme. If *name* is omitted, the text is obtained from the standard input. Omitting the `-i` option altogether creates an empty SCCS file. You can only create one SCCS file with an `admin -i` command. Creating more than one SCCS file with a single `admin` command requires that they be created empty, in which case the `-i` option should be omitted. Note that the `-i` option implies the `-n` option.

Initial release

`-r rel`

Initial release: the *rel*ease into which the initial delta is inserted. `-r` may be used only if the `-i` option is also used. The initial delta is inserted into release 1 if the `-r` option is not used. The level of the initial delta is always 1, and initial deltas are named 1.1 by default.

Descriptive text

`-t [name]`

Descriptive text: The file *name* contains descriptive text for the SCCS file. The descriptive text file name *must* be supplied when creating a new SCCS file (either or both `-n` and `-i` options) and the `-t` option is used. In the case of existing SCCS files: 1) a `-t` option without a file name removes descriptive text (if any) currently in the SCCS file, and 2) a `-t` option with a file name replaces the descriptive text currently in the SCCS file with any text in the named file.

Set a flag

`-f flag`

Set *flag*: specifies a *flag*, and, possibly, a value for the *flag*, to be placed in the SCCS file. Several `-f` options may be supplied on a single `admin` command line. *Flags* and their values appear in the `FLAGS` section after this list of options.

*Delete a flag***-d *flag***

Delete *flag* from an SCCS file. The **-d** option may be specified only when processing existing SCCS files. Several **-d** options may be supplied on a single `admin` command. See the FLAGS section below.

*Unlock releases***-l *list***

Unlock the specified *list* of releases. See the **-f** option for a description of the **l** flag and the syntax of a *list*.

*Add login name***-a *login***

Add *login* name, or numerical UNIX group ID, to the list of users who may make deltas (changes) to the SCCS file. A group ID is equivalent to specifying all *login* names common to that group ID. Several **-a** options may appear on a single `admin` command line. As many *logins*, or numerical group IDs, as desired may be on the list simultaneously. If the list of users is empty, anyone may add deltas.

*Erase login name***-e *login***

Erase *login* name, or numerical group ID, from the list of users allowed to make deltas (changes) to the SCCS file. Specifying a group ID is equivalent to specifying all *login* names common to that group ID. Several **-e** options may be used on a single `admin` command line.

*Insert Comment text***-y [*comment*]**

The *comment* text is inserted into the SCCS file as a comment for the initial delta in a manner identical to that of `delta`. If the **-y** option is omitted, a default comment line is inserted in the form:

```
date and time created yy/mm/dd hh:mm:ss by login
```

The **-y** option is valid only if the **-i** and/or **-n** options are specified (that is, a new SCCS file is being created).

*Modification list***-m [*mrlist*]**

The list of Modification Requests (MR) numbers is inserted into the SCCS file as the reason for creating the initial delta in a manner identical to `delta`. The **v** flag must be set and the MR numbers are validated if the **v** flag has a value (the name of an MR number validation program). Diagnostics are displayed if the **v** flag is not set or MR validation fails.

Check Structures of SCCS file

-h Check the structure of the SCCS file (see *sccsfile(5)*), and compare a newly computed check-sum (the sum of all the characters in the SCCS file except those in the first line) with the check-sum that is stored in the first line of the SCCS file.

The **-h** option inhibits writing on the file, so that it nullifies the effect of any other options supplied, and is, therefore, only meaningful when processing existing files.

- Recompute checksum* -z recompute the SCCS file check-sum and store it in the first line of the SCCS file (see -h, above).
- Using the -z option on a truly corrupted file may prevent future detection of the corruption.
- Flags In SCCS Files** The list below is a description of the *flags* which may appear as arguments to the -f (set flags) and -d (delete flags) options.
- Branch deltas can be created* b When set, the -b option can be used on a get command to create branch deltas.
- Highest retrievable release* c *ceil*
The highest release (ceiling) which may be retrieved by a get command for editing. The ceiling is a number less than or equal to 9999. The default value for an unspecified c flag is 9999.
- Lowest retrievable release* f *floor*
The lowest release (floor) which may be retrieved by a get command for editing. The floor is a number greater than 0 but less than 9999. The default value for an unspecified f flag is 1.
- Default delta number* d *SID*
The default delta number (ID) to be used by a get command.
- No ID keywords fatal error* i Treats the 'No id keywords (ge6)' message issued by get or delta as a fatal error. In the absence of the i flag, the message is only a warning. The message is displayed if no SCCS identification keywords (see get) are found in the text retrieved or stored in the SCCS file.
- Allow concurrent edits* j Concurrent get commands for editing may apply to the same SID of an SCCS file. This allows multiple concurrent updates to the same version of the SCCS file.
- Locked releases* l *list*
A *list* of locked releases to which deltas can no longer be made. A get -e fails when applied against one of these locked releases. The *list* has the following syntax:
- <list>* ::= *<range>* | *<list>* , *<range>*
<range> ::= RELEASE NUMBER | a
- The character a in the *list* is equivalent to specifying *all releases* for the named SCCS file.
- Create null deltas* n The delta command creates a 'null' delta in each release (if any) being skipped when a delta is made in a *new* release. For example, releases 3 and 4 are skipped when making delta 5.1 after delta 2.7. These null deltas serve as 'anchor points' so that branch deltas may be created from them later. If

the `n` flag is absent from the SCCS file, skipped releases will be non-existent in the SCCS file, preventing branch deltas from being created from them in the future.

`q text`

`text` is defined by the user. The `text` is substituted for all occurrences of the `%Q%` keyword in SCCS file text retrieved by `get`.

Module Name

`m module`

Module name of the SCCS file substituted for all occurrences of the `%M%` keyword in SCCS file text retrieved by `get`. If the `m` flag is not specified, the value assigned is the name of the SCCS file with the leading `s.` removed.

Module Type

`t type`

Type of module in the SCCS file substituted for all occurrences of `%Y%` keyword in SCCS file text retrieved by `get`.

Validity checking program

`v [program]`

Validity checking *program*: `delta` prompts for Modification Request (MR) numbers as the reason for creating a delta. The optional *program* specifies the name of an MR number validity checking program (see `delta`). If this flag is set when creating an SCCS file, the `-m` option must also be used even if its value is null.

Files Used

The last component of all SCCS file names must be of the form `s.file-name`. New SCCS files are given mode 444 (see `chmod`). Write permission in the pertinent directory is, of course, required to create a file. All writing done by `admin` is to a temporary x-file, called `x.file-name`, (see `get`), created with mode 444 if the `admin` command is creating a new SCCS file, or with the same mode as the SCCS file if it exists. After successful execution of `admin`, the SCCS file is removed (if it exists), and the x-file is renamed with the name of the SCCS file. This ensures that changes are made to the SCCS file only if no errors occurred.

It is recommended that directories containing SCCS files be mode 755 and that SCCS files themselves be mode 444. The mode of the directories allows only the owner to modify SCCS files contained in the directories. The mode of the SCCS files prevents any modification at all except by SCCS commands.

If it should be necessary to patch an SCCS file for any reason, the mode may be changed to 644 by the owner allowing use of a text editor. *Care must be taken!* The edited file should *always* be processed by an `admin -h` to check for corruption followed by an `admin -z` to generate a proper check-sum. Another `admin -h` is recommended to ensure the SCCS file is valid.

`admin` also uses a transient lock file (called `z.file-name`), to prevent simultaneous updates to the SCCS file by different users. See `get` for further information.

Examples of Using `admin`

Suppose you have a file called `lang` that contains a list of programming languages:

```
tutorial% cat lang
C
PL/I
FORTRAN
COBOL
Algol
tutorial%
```

We wish to give SCCS custody of 'lang' by using `admin` (which *administers* SCCS files) to create an SCCS file and initialize delta 1.1. To do so, we use `admin` as shown, and `admin` responds with a message:

```
tutorial% admin -ilang s.lang
No id keywords (cm7)
tutorial%
```

All SCCS files *must* have names that begin with 's.', hence, 's.lang'. The `-i` option, together with its value 'lang', indicates that `admin` is to create a new SCCS file and *initialize* it with the contents of the file 'lang'. This initial version is a set of changes applied to the null SCCS file; it is delta 1.1.

The message is a warning message (which may also be issued by other SCCS commands) that you can ignore for the present.

Remove the file 'lang' now — it can easily be reconstructed with the `get` command, described in section

Inserting Commentary for the Initial Delta

You can use the `-y` and `-m` options with `admin`, just as with `delta`, to insert initial descriptive commentary and/or MR numbers when an SCCS file is created. If you don't use `-y` to comment, `admin` automatically inserts a comment line of the form:

```
date and time created YY/MM/DD HH:MM:SS by logname
```

If you want to supply MR numbers (`-m` option), the `v` flag must also be set (using the `-f` option described below). The `v` flag simply determines whether or not MR numbers must be supplied when using any SCCS command that modifies a *delta commentary* in the SCCS file (see *sccsfile(5)*). Thus:

```
tutorial% admin -ifirst -mnum1 -fv s.abc
```

Note that the `-y` and `-m` options are only effective if a new SCCS file is being created.

Initializing and Modifying SCCS File Parameters

The portion of the SCCS file reserved for *descriptive text* may be initialized or changed through the use of the `-t` option. The descriptive text is intended as a summary of the contents and purpose of the SCCS file; actually its contents and

length are up to you.

When an SCCS file is being created and the `-t` option is supplied, it must be followed by the name of a file from which the descriptive text is to be taken. For example, the command

```
tutorial% admin -ifirst -tdesc s.abc
```

specifies that the descriptive text is to be taken from file 'desc'.

When processing an *existing* SCCS file, the `-t` option specifies that the descriptive text (if any) currently in the file is to be *replaced* with the text in the named file. Thus:

```
tutorial% admin -tdesc s.abc
```

specifies that the descriptive text of the SCCS file is to be replaced by the contents of 'desc'. Omitting the filename after the `-t` option *removes* the descriptive text from the SCCS file:

```
tutorial% admin -t s.abc
```

The *flags* — see the section entitled *Descriptive Text* — of an SCCS file may be initialized and changed with the `-f` (flag) option, or may be deleted with the `-d` (delete) option. The flags of an SCCS file direct certain actions of the various commands. See `admin` for a description of all the flags. For example, the `i` flag specifies that the warning message stating there are no ID keywords contained in the SCCS file should be treated as an error, and the `d` (default SID) flag specifies the default version of the SCCS file to be retrieved by the `get` command. The `-f` option sets a flag and, possibly, sets its value. For example:

```
tutorial% admin -ifirst -fi -fmodname s.abc
```

sets the `i` flag and the `m` (module name) flag. The value 'modname' specified for the `m` flag is the value that the `get` command uses to replace the `%M%` ID keyword. In the absence of the `m` flag, the name of the *g-file* is used as the replacement for the `%M%` ID keyword. Note that several `-f` options may be supplied on a single `admin` command, and that `-f` options may be supplied whether the command is creating a new SCCS file or processing an existing one.

The `-d` option deletes a flag from an SCCS file, and may only be specified when processing an existing file. As an example, the command:

```
tutorial% admin -dm s.abc
```

removes the `m` flag from the SCCS file. Several `-d` options may be supplied on a single `admin` command, and may be interspersed with `-f` options.

SCCS files contain a list (*user list*) of login names and/or group IDs of users who are allowed to create deltas. This list is normally empty, implying that *anyone* may create deltas. To add login names and/or group IDs to the list, use the `admin` command with the `-a` option. For example:

```
tutorial% admin -awendy -aalison -a1234 s.abc
```

adds the login names 'wendy' and 'alison' and the group ID '1234' to the list. The `-a` option may be used whether `admin` is creating a new SCCS file or processing an existing one, and may appear several times. The `-e` option is used in an analogous manner if one wishes to remove ('erase') login names or group IDs from the list. A.9.

A.6. `cdc` — Change Delta Commentary

`cdc` changes the *delta commentary*, for the SID specified by the `-r` option, of each named SCCS file.

```
cdc -rSID [-m[mrlist] ] [ -y [comment] ] filename . . .
```

Delta commentary is defined to be the Modification Request (MR) and comment information normally specified via the `delta` command (`-m` and `-y` options).

If a directory is named, `cdc` behaves as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the path name does not begin with `s.`) and unreadable files are silently ignored. If a name of `-` is given, the standard input is read (see the NOTES below) each line of the standard input is taken to be the name of an SCCS file to be processed.

Arguments to `cdc`, which may appear in any order, consist of options and file names.

`cdc` Options

All the described options apply independently to each named file:

ID String

`-rSID`

Specifies the *SCCS ID* entification string of a delta for which the delta commentary is to be changed.

MR List

`-m[mrlist]`

If the SCCS file has the `v` flag set (see `admin`), a list of MR numbers to be added and/or deleted in the delta commentary of the SID specified by the `-r` option *may* be supplied. A null MR list has no effect.

MR entries are added to the list of MRs in the same manner as that of `delta`. To delete an MR, precede the MR number with the character `!` (see EXAMPLES). If the MR to be deleted is currently in the list of MRs, it is removed and changed into a "comment" line. A list of all deleted MRs is placed in the comment section of the delta commentary and preceded by a comment line stating that they were deleted.

If `-m` is not used and the standard input is a terminal, the prompt `MRs?` is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. The `MRs?` prompt always precedes the `comments?` prompt (see `-y` option).

MRs in a list are separated by blanks and/or tab characters. An unescaped new-line character terminates the MR list.

Note that if the `v` flag has a value (see `admin`), it is taken to be the name of a program (or shell procedure) which validates the correctness of the MR numbers. If a non-zero exit status is returned from the MR number validation program, `cdc` terminates and the delta commentary remains unchanged.

Comment text

`-y[comment]`

Arbitrary text used to replace the *comment(s)* already existing for the delta specified by the `-r` option. The previous comments are kept and preceded by a comment line stating that they were changed. A null *comment* has no effect.

If `-y` is not specified and the standard input is a terminal, the prompt `comments?` is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. An unescaped new-line character terminates the *comment* text.

Examples of Using `cdc`

```
tutorial% cdc -r1.6 -m"b178-12345 !b177-54321 b179-00001" -ytrouble s.file
```

adds `b178-12345` and `b179-00001` to the MR list, removes `b177-54321` from the MR list, and adds the comment `trouble` to delta 1.6 of `s.file`.

```
tutorial% cdc -r1.6 s.file
MRs? !b177-54321 b178-12345 b179-00001
comments? trouble
```

does the same thing.

NOTE *If SCCS file names are supplied to the `cdc` command via the standard input (`-` on the command line), then the `-m` and `-y` options must also be used.*

Files Used

x-file (see `delta`)
z-file (see `delta`)

A.7. `comb` — Combine SCCS Deltas

`comb` generates a Bourne Shell procedure which, when run, will reconstruct the given SCCS files.

```
comb [ -o ] [ -s ] [ -pSID ] [ -c list ] filename . . .
```

If a directory is named, `comb` behaves as though each file in the directory were

specified as a named file, except that non-SCCS files (last component of the path name does not begin with `s.`) and unreadable files are silently ignored. If a name of `-` is given, the standard input is read; each line of the standard input is taken to be the name of an SCCS file to be processed; non-SCCS files and unreadable files are silently ignored. The generated shell procedure is written on the standard output.

`comb` Options

Options are explained as though only one named file is to be processed, but the effects of any option apply independently to each named file.

ID String

`-p SID`

The *SCCS IDentification* string (SID) of the oldest delta to be preserved. All older deltas are discarded in the reconstructed file.

Preserve list

`-c list`

A *list* of deltas to be preserved. All other deltas are discarded. See `get` for the syntax of a *list*.

Access at release

`-o` For each `get -e` generated, the reconstructed file is accessed at the release of the delta to be created. In the absence of the `-o` option, the reconstructed file is accessed at the most recent ancestor. Use of the `-o` option may decrease the size of the reconstructed SCCS file. It may also alter the shape of the delta tree of the original file.

Generate report

`-s` Generate a shell procedure which, when run, will produce a report giving, for each file: the file name, size (in blocks) after combining, original size (also in blocks), and percentage change computed by:

$$100 * (\text{original combined}) / \text{original}$$

It is recommended that before any SCCS files are actually combined, you should use this option to determine exactly how much space is saved by the combining process.

If no options are specified, `comb` preserves only leaf deltas and the minimal number of ancestors needed to preserve the tree.

Files Used

`s.COMB`

The name of the reconstructed SCCS file.

`comb?????`

Temporary.

Limitations of the `comb` Command

`comb` may rearrange the shape of the tree of deltas. It may not save any space; in fact, it is possible for the reconstructed file to actually be larger than the original.

A.8. delta — Make a Delta `delta` permanently introduces into the named SCCS file changes that were made to the file retrieved by `get` (called the *g-file*, or generated file).

```
delta [ -rSID ] [ -s ] [ -n ] [ -glist ] [ -m [ mrlist ] ] [ -y [ comment ] ] [ -p ] filename . . .
```

`delta` makes a delta to each named SCCS file. If a directory is named, `delta` behaves as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the path name does not begin with `s.`) and unreadable files are silently ignored. If a name of `-` is given, the standard input is read (see WARNINGS; each line of the standard input is taken to be the name of an SCCS file to be processed).

`delta` may issue prompts on the standard output depending upon certain options specified and flags (see `admin`) that may be present in the SCCS file (see `-m` and `-y` options below).

delta Options

Options apply independently to each named file.

Delta number

`-r SID`

Uniquely identifies which delta is to be made to the SCCS file. The use of this option is necessary only if two or more outstanding `get`'s for editing (`get -e`) on the same SCCS file were done by the same person (login name). The SID value specified with the `-r` option can be either the SID specified on the `get` command line or the SID to be made as reported by the `get` command (see `get`). A diagnostic results if the specified SID is ambiguous, or, if necessary and omitted on the command line.

No report

`-s` Do not display the created delta's ID, number of lines inserted, deleted and unchanged in the SCCS file.

Retain g-file

`-n` Retain the edited *g-file* which is normally removed at completion of delta processing.

Ignore list

`-g list`

Specifies a *list* of deltas to be *ignored* when the file is accessed at the change level (ID) created by this delta. See `get` for the definition of *list*.

MR number

`-m [mrlist]`

If the SCCS file has the `v` flag set (see `admin`), a Modification Request (MR) number *must* be supplied as the reason for creating the new delta.

If `-m` is not used and the standard input is a terminal, the prompt `MRs?` is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. The `MRs?` prompt always precedes the `comments?` prompt (see `-y` option).

MR's in a list are separated by blanks and/or tab characters. An unescaped new-line character terminates the MR list.

Note that if the `v` flag has a value (see `admin`), it is taken to be the name of a program (or shell procedure) which will validate the correctness of the MR numbers. If a non-zero exit status is returned from MR number validation program, `delta` terminates (it is assumed that the MR numbers were not all valid).

Comment text`-y [comment]`

Arbitrary text to describe the reason for making the delta. A null string is considered a valid *comment*.

If `-y` is not specified and the standard input is a terminal, the prompt `comments?` is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. An unescaped new-line character terminates the comment text.

Display differences`-p` Display (on the standard output) the SCCS file differences before and after the delta is applied in a `diff` format.

Files Used

g-file Existed before the execution of `delta`; removed after completion of `delta`.

p-file Existed before the execution of `delta`; may exist after completion of `delta`.

q-file Created during the execution of `delta`; removed after completion of `delta`.

x-file Created during the execution of `delta`; renamed to SCCS file after completion of `delta`.

z-file Created during the execution of `delta`; removed during the execution of `delta`.

d-file Created during the execution of `delta`; removed after completion of `delta`.

`/bin/diff`
Program to compute differences between the “gotten” file and the *g-file*.

NOTE *Lines beginning with an ASCII SOH character (binary 001) cannot be placed in the SCCS file unless the SOH is escaped. This character has special meaning to SCCS (see `sccsfile(5)`) and will cause an error.*

NOTE *A `get` of many SCCS files, followed by a `delta` of those files, should be avoided when the `get` generates a large amount of data. Instead, multiple `get/delta` sequences should be used.*

NOTE *If the standard input (`-`) is specified on the `delta` command line, the `-m` (if necessary) and `-y` options must also be present. Omission of these options is an error.*

Examples of Using `delta`

To record the changes that were applied to 'lang' within the SCCS file, use the `delta` command. `delta` asks for comments describing the change, and you respond with a description of why the changes were made:

```
tutorial% delta s.lang
comments? added SNOBOL and Ratfor
      More messages from delta — see below
tutorial%
```

`delta` then reads the *p-file* and determines what changes were made to the file 'lang'. `delta` does this by doing its own `get` to retrieve the original version, and then applying `diff(1)` to the original version and the edited version. When the changes to 'lang' have been stored in 's.lang', the dialogue with `delta` looks like:

```
tutorial% delta s.lang
comments? added SNOBOL and Ratfor
1.2
2 inserted
0 deleted
5 unchanged
tutorial%
```

The number '1.2' is the name of the delta just created, and the next three lines are a summary of the changes made to 's.lang'.

More Notes on `delta`

`delta` does a series of checks before creating the delta:

1. Searches the *p-file* for an entry containing the user's login name, because the user who retrieved the *g-file* must be the one who creates the delta. `delta` displays an error message if the entry is not found. Note that if the login name of the user appears in more than one entry (that is, the same user did a `get -e` more than once on the same SCCS file), the `-r` option must be used with `delta` to specify an SID that uniquely identifies the *p-file* entry¹.
2. Performs the same permission checks as `get -e`.

If these checks succeed, `delta` compares the *g-file* (via `diff(1)`) with its own, temporary copy of the *g-file* as it was before editing, to determine what has been changed. This temporary copy of the *g-file* is called the *d-file* (its name is formed by replacing the 's.' of the SCCS file name with 'd.');

`delta` retrieves it by doing its own `get` at the SID specified in the *p-file* entry. If you would like to see the results of `delta`'s *diff*, use the `-p` option to display it on standard output.

In practice, the most common use of `delta` is:

¹ The SID specified may be either the SID retrieved by `get`, or the SID `delta` is to create.

```
tutorial% delta s.abc
```

If your standard output is a terminal, `delta` replies: ‘comments?’. You may now type a response — usually a description of why the delta is being made — of up to 512 characters, terminating with a newline character. Newline characters *not* intended to terminate the response should be preceded by ‘\’.

If the SCCS file has a `v` flag, `delta` asks for ‘MRs?’ before prompting for ‘comments?’ (again, this prompt is printed only if the standard output is a terminal). Enter MR² numbers, separated by blanks and/or tabs, and terminate your response with a newline character.

If you want to enter commentary (comments and/or MR numbers) directly on the command line, use the `-y` and/or `-m` options, respectively. For example:

```
tutorial% delta -y"descriptive comment" -m"mrnum1 mrnum2" s.abc
```

inserts the ‘descriptive comment’ and the MR numbers ‘mrnum1’ and ‘mrnum2’ without prompting or reading from standard input. `-m` can only be used if the SCCS file has a `v` flag. These options are useful when `delta` is executed from within a Shell procedure.

The commentary (comments and/or MR numbers), whether solicited by `delta` or supplied via options, is recorded as part of the entry for the delta being created, and applies to *all* SCCS files processed by the same invocation of `delta`. Thus if `delta` is used with more than one file argument, and the first file named has a `v` flag, all files named must have this flag. Similarly, if the first file named does not have this flag, then none of the files named may have it. Only files conforming to these rules are processed.

After the prompts for commentary, and before any other output, `delta` displays:

```
No id keywords (cm7)
```

if it finds no ID keywords in the edited *g-file* while making a delta. If there *were* any ID keywords in the SCCS file, this might mean one of two things. The keywords may have been replaced by their values (if a `get` without the `-e` option was used to retrieve the *g-file*). Or, the keywords may have been accidentally deleted or changed while editing the *g-file*. Of course, the file may never have had any ID keywords. In any case, it is left up to you to decide whether any action is necessary, but the delta is made regardless (unless there is an `i` flag in the SCCS file, which makes this a fatal error and kills the delta).

When processing is complete, `delta` displays a message containing the SID of the created delta (obtained from the *p-file* entry), and the counts of lines inserted, deleted, and left unchanged. Thus, a typical message might be:

² In a tightly controlled environment, one would expect deltas to be created only as a result of some trouble report, change request, trouble ticket, etc. (collectively called here Modification Requests, or MRs) and would think it desirable or necessary to record such MR number(s) within each delta.

```

1.4
14 inserted
7 deleted
345 unchanged

```

The reported counts may not agree with your sense of changes made; there are a number of ways to describe a set of such changes, especially if lines are moved around in the *g-file*, and `delta` may describe the set differently than you. 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*.

After processing of an SCCS file is complete, the corresponding *p-file* entry is removed from the *p-file*³. If there is only *one* entry in the *p-file*, the *p-file* itself is removed.

In addition, `delta` removes the edited *g-file*, unless the `-n` option is specified. Thus:

```
tutorial% delta -n s.abc
```

keeps the *g-file* upon completion of processing.

The `-s` (silent) option suppresses all output that is normally directed to the standard output, except the initial prompts for commentary. If you use `-s` with `-y` (and, possibly, `-m`), `delta` neither reads standard input nor writes to standard output.

A.9. `get` — Get Version of SCCS File

`get` generates an ASCII text file from each named SCCS file according to the specified option. Arguments may be specified in any order, options apply to all named SCCS files. If a directory is named, `get` behaves as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the path name does not begin with `s.`) and unreadable files are silently ignored. If a name of `-` is given, the standard input is read; each line of the standard input is taken to be the name of an SCCS file to be processed. Again, non-SCCS files and unreadable files are silently ignored.

```
get [ -rSID ] [ -ccutoff ] [ -ilist ] [ -xlist ] [ -aseq-no. ] [ -k ] [ -e ]
    [ -l[p] ] [ -p ] [ -m ] [ -n ] [ -s ] [ -b ] [ -g ] [ -t ] filename ...
```

The generated text is normally written into a file called the *g-file* whose name is derived from the SCCS file name by simply removing the leading `s.`; (see also *FILES*, below).

³ All 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* described above.

- get Options** Options are explained below as though only one SCCS file is to be processed, but the effects of any option argument applies independently to each named file.
- ID string* **-r *SID***
The string (ID) of the version (delta) of an SCCS file to be retrieved. Table 1 below shows, for the most useful cases, what version of an SCCS file is retrieved (as well as the ID of the version to be eventually created by `delta` if the `-e` option is also used), as a function of the *SID* specified.
- Cutoff* **-c *cutoff***
Cutoff date-time, in the form: YY[MM[DD[HH[MM[SS]]]]]
No changes (deltas) to the SCCS file which were created after the specified *cutoff* date-time are included in the generated ASCII text file. Units omitted from the date-time default to their maximum possible values; that is, `-c7502` is equivalent to `-c750228235959`. Any number of non-numeric characters may separate the various 2 digit pieces of the *cutoff* date-time. This feature allows one to specify a *cutoff* date in the form: `-c77/2/29:22:25`. Note that this implies that one may use the `%E%` and `%U%` identification keywords.
- Get for editing* **-e** This `get` is for editing or making a change (delta) to the SCCS file via a subsequent use of `delta`. A `get -e` applied to a particular version (ID) of the SCCS file prevents further `get -e` commands on the same *SID* until `delta` is run or the `j` (joint edit) flag is set in the SCCS file (see `admin`). Concurrent use of `get -e` for different IDs is always allowed.
If the *g-file* generated by a `get -e` is accidentally ruined in the process of editing it, it may be regenerated by re-running a `get` with the `-k` option in place of the `-e` option.
SCCS file protection specified via the ceiling, floor, and authorized user list stored in the SCCS file (see `admin`) are enforced when the `-e` option is used.
- New branch* **-b** Used with the `-e` option to indicate that the new delta should have an *SID* in a new branch as shown in Table 1. This option is ignored if the `b` flag is not present in the file (see `admin`) or if the retrieved `delta` is not a leaf *delta*. A leaf *delta* is one that has no successors on the SCCS file tree.
- NOTE *A branch delta may always be created from a non-leaf delta.*
- Include list* **-i *list***
A *list* of deltas to be included (forced to be applied) in the creation of the generated file. The *list* has the following syntax:
`<list> ::= <range> | <list> , <range>`
`<range> ::= ID | ID-ID`
ID, the SCCS Identification of a delta, may be in any form shown in the 'ID Specified' column of Table 1. Partial IDs are interpreted as shown in the 'ID

Retrieved' column of Table 1.

- Exclude list* -x *list*
A *list* of deltas to be excluded (forced not to be applied) in the creation of the generated file. See the -i option for the *list* format.
- Don't expand ID keywords* -k Do not replace identification keywords (see below) in the retrieved text by their value. The -k option is implied by the -e option.
- Write delta summary* -l [*p*]
Write a delta summary into an *l-file*. If -lp is used, the delta summary is written on the standard output and the *l-file* is not created. See *FILES* for the format of the *l-file*.
- Write text to standard output* -p Write the text retrieved from the SCCS file to the standard output. No *g-file* is created. All output which normally goes to the standard output goes to the standard error file instead, unless the -s option is used, in which case it disappears.
- Suppress all output* -s Suppress all output normally written on the standard output. However, fatal error messages (which always go to the standard error file) remain unaffected.
- Show delta IDs* -m Precede each text line retrieved from the SCCS file with the ID of the delta that inserted the text line in the SCCS file. The format is: ID, followed by a horizontal tab, followed by the text line.
- Show Module names* -n Precede each generated text line with the %M% identification keyword value (see below). The format is: %M% value, followed by a horizontal tab, followed by the text line. When both the -m and -n options are used, the format is: %M% value, followed by a horizontal tab, followed by the -m option generated format.
- Don't retrieve text* -g Do not actually retrieve text from the SCCS file. It is primarily used to generate an *l-file*, or to verify the existence of a particular ID.
- Access top delta* -t Access the most recently created ('top') delta in a given release (for example, -r1), or release and level (for example, -r1.2).
- Delta sequence number* -a *seq-no.*
The delta sequence number of the SCCS file delta (version) to be retrieved (see *sccsfile* (5)). This option is used by the *comb* command; it is not a generally useful option, and users should not use it. If both the -r and -a options are specified, the -a option is used. Care should be taken when using the -a option in conjunction with the -e option, as the SID of the delta to be created may not be what one expects. The -r option can be used with the -a and -e options to control the naming of the SID of the delta to be created.

For each file processed, `get` responds (on the standard output) with the SID being accessed and with the number of lines retrieved from the SCCS file.

If the `-e` option is used, the SID of the delta to be made appears after the SID accessed and before the number of lines generated. If there is more than one named file or if a directory or standard input is named, each file name is printed (preceded by a new-line) before it is processed. If the `-i` option is used included deltas are listed following the notation 'Included'; if the `-x` option is used, excluded deltas are listed following the notation 'Excluded'.

Table A-1 *Determination of SCCS Identification String*

<i>SID*</i> <i>Specified</i>	<i>-b Option</i> <i>Used†</i>	<i>Other</i> <i>Conditions</i>	<i>SID</i> <i>Retrieved</i>	<i>SID of Delta</i> <i>to be Created</i>
none‡	no	R defaults to mR	mR.mL	mR.(mL+1)
none‡	yes	R defaults to mR	mR.mL	mR.mL.(mB+1).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 R does <i>not</i> exist	hR.mL**	hR.mL.(mB+1).1
R	—	Trunk succ.# in release > R and R exists	R.mL	R.mL.(mB+1).1
R.L	no	No trunk succ.	R.L	R.(L+1)
R.L	yes	No trunk succ.	R.L	R.L.(mB+1).1
R.L	—	Trunk succ. in release ≥ R	R.L	R.L.(mB+1).1
R.L.B	no	No branch succ.	R.L.B.mS	R.L.B.(mS+1)
R.L.B	yes	No branch succ.	R.L.B.mS	R.L.(mB+1).1
R.L.B.S	no	No branch succ.	R.L.B.S	R.L.B.(S+1)
R.L.B.S	yes	No branch succ.	R.L.B.S	R.L.(mB+1).1
R.L.B.S	—	Branch succ.	R.L.B.S	R.L.(mB+1).1

* 'R', 'L', 'B', and 'S' are the 'release', 'level', 'branch', and 'sequence' components of the SID, respectively; '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 *new* branch (that is, maximum branch number plus one) of level L within release R'. 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.

- ** 'hR' is the highest *existing* release that is lower than the specified, *nonexistent*, release R.
- ***
Forces creation of the *first* delta in a *new* release.
- # Successor.
- † The `-b` option is effective only if the `b` flag (see `admin`) is present in the file. An entry of `-` means 'irrelevant'.
- ‡ This case 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 interpreted as if it had been specified on the command line. Thus, one of the other cases in this table applies.

Identification Keywords

When you generate a *g-file* to be used for compilation, it is useful and informative to record the date and time of creation, the version retrieved, the module's name, etc., within the *g-file*, so that this information appears in a load module when one is eventually created. SCCS provides a convenient mechanism for doing this automatically. *Identification (ID) keywords* appearing anywhere in the generated file are replaced by appropriate values according to the definitions of these ID keywords.

The format of an ID keyword is an upper-case letter enclosed by percent signs (%). For example, %I% is an ID keyword that is replaced by the SID of the retrieved version of a file. Similarly, %H% is an ID keyword for the current date (in the form 'mm/dd/yy'), and %M% is the name of the *g-file*.

Thus, using `get` on an SCCS file that contains the C declaration:

```
char identification [ ] = "%M% %I% %H%";
```

gives (for example) the following:

```
char identification [ ] = "modulename 2.3 03/17/83";
```

If there are no ID keywords in the text, `get` might display:

```
No id keywords (cm7)
tutorial%
```

This message is normally treated as a warning by `get`. However, if an `i` flag is present in the SCCS file, it is treated as an error — see section A.8 for further information.

Table A-2 *Identification Keywords*

<i>Keyword</i>	<i>Value</i>
%M%	Module name: either the value of the <code>m</code> flag in the file (see <code>admin</code>), or if absent, the name of the SCCS file with the leading <code>s.</code> removed.
%I%	SCCS identification (ID) (%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: value of the <code>t</code> flag in the SCCS file (see <code>admin</code>).
%F%	SCCS file name.
%P%	Fully qualified SCCS file name.
%Q%	The value of the <code>q</code> flag in the file (see <code>admin</code>).
%C%	Current line number. This keyword is intended for identifying messages output by the program such as 'this shouldn't have happened' type errors. It is <i>not</i> intended to be used on every line to provide sequence numbers.
%Z%	The 4-character string @ (#) recognizable by what.
%W%	A shorthand notation for constructing what strings for UNIX program files. %W% = %Z%%M%<tab>%I%
%A%	Another shorthand notation for constructing what strings for non-UNIX program files. %A% = %Z%%Y% %M% %I%%Z%

Retrieving Different Versions

You can retrieve versions other than the default version of an SCCS file by using various options. Normally, the default version is the most recent delta of the highest-numbered release on the *trunk* of the SCCS file tree. However, if the SCCS file being processed has a `d` (default SID) flag, the SID specified as the value of this flag is used as a default. The default SID is interpreted in exactly the same way as the value supplied with the `-r` option of `get`.

The `-r` option specifies an SID to be retrieved, in which case the `d` (default SID) flag (if any) is ignored. For example, to retrieve version 1.3 of file 's.abc', type:

```
tutorial% get -r1.3 s.abc
1.3
64 lines
tutorial%
```

A branch delta may be retrieved in the same way:

```
tutorial% get -r1.5.2.3 s.abc
1.5.2.3
234 lines
tutorial%
```

When a two- or four-component SID is specified as a value for the `-r` option (as above) and the particular version does not exist in the SCCS file, an error message results.

If you omit the level number of the SID, `get` retrieves the *trunk* delta with the highest level number within the given release, if the given release exists:

```
tutorial% get -r3 s.abc
3.7
213 lines
tutorial%
```

`get` retrieved delta 3.7, the highest level trunk delta in release 3. If the given release does not exist, `get` goes to the next-highest existing release, and retrieves the *trunk* delta with the highest level number. For example, if release 9 does not exist in file 's.abc', and release 7 is actually the highest-numbered release below 9, then `get` would generate:

```
tutorial% get -r9 s.abc
7.6
420 lines
tutorial%
```

indicating that trunk delta 7.6 is the latest version of file 's.abc' below release 9.

Similarly, if you omit the sequence number of an SID, as in:

```
tutorial% get -r4.3.2 s.abc
4.3.2.8
89 lines
tutorial%
```

`get` retrieves the branch delta with the highest sequence number on the given branch, if it exists. If the given branch does not exist, an error message results.

The `-t` option retrieves the latest ('top') version in a particular *release* (that is, when no `-r` option is supplied, or when its value is simply a release number). The latest version is defined as that delta which was produced most recently, independent of its location on the SCCS file tree. Thus, if the most recent delta in release 3 is trunk delta 3.5, doing a `get -t` on release 3 produces:

```
tutorial% get -r3 -t s.abc
3.5
59 lines
tutorial%
```

However, if branch delta 3.2.1.5 were the latest delta (created after delta 3.5), the same command produces:

```
tutorial% get -r3 -t s.abc
3.2.1.5
46 lines
tutorial%
```

Retrieving to Make Changes

Specifying the `-e` option to the `get` command indicates the intent to make a delta sometime later, and, as such, its use is restricted. If the `-e` option is present, `get` checks the following things:

1. The *user list*, the list of *login* names and/or *group IDs* of users allowed to make deltas, to determine if the login name or group ID of the user executing `get` is on that list. Note that a *null* (empty) user list behaves as if it contained *all* possible login names.
2. That the *release* (R) of the version being retrieved satisfies the relation:

$$\text{floor} \leq R \leq \text{ceiling}$$
 to determine if the release being accessed is a protected release. The *floor* and *ceiling* are specified as *flags* in the SCCS file.
3. That the *release* (R) is not *locked* against editing. The *lock* is specified as a flag in the SCCS file.
4. Whether or not *multiple concurrent edits* are allowed for the SCCS file as specified by the `j` flag in the SCCS file. Multiple concurrent edits are described in the section entitled *Concurrent Edits of the Same SID*.

`get` terminates processing of the corresponding SCCS file if any of the first three conditions fails.

If the above checks succeed, `get` with the `-e` option creates a *g-file* in the current directory with mode 644 (readable by everyone, writable only by the owner) owned by the real user.

`get` terminates with an error if a *writable g-file* already exists — 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` when the `-e` option is specified, because the generated *g-file* is to be subsequently used to create another delta, and replacement of ID keywords would permanently change them within the SCCS file. In view of this, `get` does not check for the presence of ID keywords within the *g-file*, so that the message: ‘No id keywords (cm7)’ is

never displayed when `get` is invoked with the `-e` option.

In addition, a `get` with the `-e` option creates (or updates) a *p-file*, for passing information to the `delta` command. Let's look at an example of `get -e`:

```
tutorial% get -e s.abc
1.3
new delta 1.4
67 lines
tutorial%
```

The message indicates that `get` has retrieved version 1.3, which has 67 lines; the version `delta` will create is version 1.4.

If the `-r` and/or `-t` options are used together with the `-e` option, the version retrieved for editing is as specified by the `-r` and/or `-t` options.

The options `-i` and `-x` may be used to specify a list of deltas to be *included* and *excluded*, respectively, by `get`. See `get` for the syntax of such a list. 'Including a delta' forces the changes that constitute the particular delta to be included in the retrieved version — this is useful for applying the same changes to more than one version of the SCCS file. 'Excluding a delta' forces it *not* to be applied. This is useful for undoing the effects of a previous delta in the version of the SCCS file to be created.

Whenever deltas are included or excluded, `get` checks for possible interference between such deltas and those 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 displays the range of lines within the retrieved *g-file* in which the problem may exist. The user is expected to examine the *g-file* to determine whether a problem actually exists, and to take whatever corrective measures are deemed necessary.

NOTE *The `-i` and `-x` options should be used with extreme care.*

The `-k` option to `get` can be used to regenerate a *g-file* that may have been accidentally removed or ruined after executing `get` with the `-e` option, or to simply generate a *g-file* in which the replacement of ID keywords has been suppressed. Thus, a *g-file* generated by the `-k` option is identical to one produced by `get` executed with the `-e` option. However, no processing related to the *p-file* takes place.

Concurrent Edits of Different SIDs

The ability to retrieve different versions of an SCCS file allows a number of deltas to be 'in progress' at any given time. In general, several people may simultaneously edit the same SCCS file provided they are editing *different versions* of that file. This is the situation we discuss in this section. However, there is a provision for multiple concurrent edits, so that more than one person can edit the *same version* — see the section entitled *Concurrent Edits of the Same SID*.

The *p-file* — created via a `get -e` command — is named by replacing the 's.' in the SCCS file name with 'p.'. The *p-file* is created in the directory containing

the SCCS file, is given mode 644 (readable by everyone, writable only by the owner), and is owned by the effective user. The *p-file* contains the following information for each delta that is still ‘in progress’:⁴

- The SID of the retrieved version.
- The SID that will be given to the new delta when it is created.
- The login name of the real user executing `get`.

The first execution of `get -e` creates the *p-file* for the corresponding SCCS file. Subsequent executions only *update* the *p-file* by inserting a line containing the above information. Before inserting this line, however, `get` performs two checks. First, it searches the entries in the *p-file* for an SID which matches that of the requested version, to make sure that the requested version has not already been retrieved. Secondly, `get` determines whether or not multiple concurrent edits are allowed. If the requested version has been retrieved and multiple concurrent edits are not allowed, an error message results. Otherwise, the user is informed that other deltas are in progress, and processing continues.

It is important to note that the various executions of `get` should be carried out from different directories. Otherwise, only the first use of `get` will succeed; since subsequent `gets` would attempt to overwrite a *writable g-file*, they produce an SCCS error condition. In practice, this problem does not arise: normally such multiple executions are performed by different users⁵ from different working directories.

Table A-1 shows, for the most useful cases, what version of an SCCS file is retrieved by `get`, as well as the SID of the version to be eventually created by `delta`, as a function of the SID specified to `get`.

Concurrent Edits of the Same SID

Normally, `gets` for editing (`-e` option specified) cannot operate concurrently on the same SID. Usually `delta` must be used before another `get -e` on the same SID. However, multiple concurrent edits (two or more *successive* `get -e` commands based on the same retrieved SID) *are* allowed if the `j` flag is set in the SCCS file. Thus:

```
tutorial% get -e s.abc
1.1
new delta 1.2
5 lines
tutorial%
```

may be immediately followed by:

⁴ Other information may be present, but is not of concern here. See `get` for further discussion.

⁵ See the section entitled *Protection* for a discussion of how different users can use SCCS commands on the same files.

```
tutorial% get -e s.abc
1.1
new delta 1.1.1.1
5 lines
tutorial%
```

without an intervening use of `delta`. In this case, a `delta` command corresponding to the first `get` produces delta 1.2 (assuming 1.1 is the latest (most recent) trunk delta), and the `delta` command corresponding to the second `get` produces delta 1.1.1.1.

Options That Affect Output

When the `-p` option is specified, `get` writes the retrieved text to the standard output, rather than to a *g-file*. In addition, 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 diagnostic output. This may be used, for example, to create *g-files* with arbitrary names:

```
tutorial% get -p s.abc > arbitrary-filename
```

The `-s` option suppresses all output that is *normally* directed to the standard output. Thus, the SID of the retrieved version, the number of lines retrieved, and so on, do not appear on the standard output. `-s` does not affect messages directed to the diagnostic output. `-s` is often used in conjunction with the `-p` option to 'pipe' the output of `get`, as in:

```
tutorial% get -p -s s.abc | nroff
```

A `get -g` verifies the existence of a particular SID in an SCCS file but does not actually retrieve the text. This may be useful in a number of ways. For example,

```
tutorial% get -g -r4.3 s.abc
```

displays the specified SID if it exists in the SCCS file, and generates an error message if it doesn't. `-g` can also be used to regenerate a *p-file* that has been destroyed:

```
tutorial% get -e -g s.abc
```

`get` used with the `-l` option creates an *l-file*, which is named by replacing the 's.' of the SCCS file name with 'l.'. This file is created in the current directory, with mode 444 (read-only), and is owned by the real user. It contains a table (format described in `get`) showing which deltas were used in constructing a particular version of the SCCS file. For example:

```
tutorial% get -r2.3 -l s.abc
```

generates an *l-file* showing which deltas were applied to retrieve version 2.3 of the SCCS file. Specifying a *value* of 'p' with the `-l` option, as in:

```
tutorial% get -lp -r2.3 s.abc
```

sends the generated output to the standard output rather than to the *l-file*. Note that the `-g` option may be used with the `-l` option to suppress the actual text retrieval.

The `-m` option identifies the origin of each change applied to an SCCS file. `-m` tags each line of the generated *g-file* with the SID of the delta it came from. The SID precedes the line, and is separated from the text by a tab character.

When the `-n` option is specified, each line of the generated *g-file* is preceded by the value of the `%M%` ID keyword and a tab character. The `-n` option is most often used in a pipeline with `grep(1)`. For example, to find all lines that match a given pattern in the latest version of each SCCS file in a directory:

```
tutorial% get -p -n -s directory | grep pattern
```

If both the `-m` and `-n` options are specified, each line of the generated *g-file* is preceded by the value of the `%M%` ID keyword and a tab (the effect of the `-n` option), followed by the line in the format produced by the `-m` option.

Since using the `-m` option, the `-n` option, or both, modifies the contents of the *g-file*, such a *g-file* must *not* be used for creating a delta. Therefore, neither the `-m` nor the `-n` option may be used with the `-e` option.

Files Used

Several auxiliary files may be created by *get*. These files are known generically as the *g-file*, *l-file*, *p-file*, and *z-file*. The letter before the hyphen is called the tag. An auxiliary file name is formed from the SCCS file name: the last component of all SCCS file names must be of the form `s.module-name`, the auxiliary files are named by replacing the leading `s` with the tag. The *g-file* is an exception to this scheme: the *g-file* is named by removing the `s.` prefix. For example, `s.xyz.c`, the auxiliary file names would be `xyz.c`, `l.xyz.c`, `p.xyz.c`, and `z.xyz.c`, respectively.

g-file

The *g-file*, which contains the generated text, is created in the current directory (unless the `-p` option is used). A *g-file* is created in all cases, whether or not any lines of text were generated by the *get*. It is owned by the real user. If the `-k` option is used or implied its mode is 644; otherwise its mode is 444. Only the real user need have write permission in the current directory.

l-file

The *l-file* contains a table showing which deltas were applied in generating the retrieved text. The *l-file* is created in the current directory if the `-l` option is used; its mode is 444 and it is owned by the real user. Only the real user need have write permission in the current directory.

Format of Lines in the *l*-file

Lines in the *l*-file have the following format:

- a. A blank character if the delta was applied; * otherwise.
- b. A blank character if the delta was applied or wasn't applied and ignored; * if the delta wasn't applied and wasn't ignored.
- c. A code indicating a 'special' reason why the delta was or was not applied:
'I': Included.
'X': Excluded.
'C': Cut off (by a `-c` option).
- d. Blank.
- e. SCCS identification (ID).
- f. Tab character.
- g. Date and time (in the form YY/MM/DD HH:MM:SS) of creation.
- h. Blank.
- i. Login name of person who created `delta`.

The comments and MR data follow on subsequent lines, indented one horizontal tab character. A blank line terminates each entry.

p-file

The *p*-file passes information resulting from a `get -e` along to `delta`. Its contents are also used to prevent a subsequent execution of a `get -e` for the same SID until `delta` is executed or the joint edit flag, `j`, (see `admin`) is set in the SCCS file. The *p*-file is created in the directory containing the SCCS file and the effective user must have write permission in that directory. Its mode is 644 and it is owned by the effective user. The format of the *p*-file is: the gotten ID, followed by a blank, followed by the SID that the new delta will have when it is made, followed by a blank, followed by the login name of the real user, followed by a blank, followed by the date-time the `get` was executed, followed by a blank and the `-i` option if it was present, followed by a blank and the `-x` option if it was present, followed by a new-line. There can be an arbitrary number of lines in the *p*-file at any time; no two lines can have the same new delta ID.

z-file

The *z*-file serves as a *lock-out* mechanism against simultaneous updates. Its contents are the binary (2 bytes) process ID of the command (that is, `get`) that created it. The *z*-file is created in the directory containing the SCCS file for the duration of `get`. The same protection restrictions as those for the *p*-file apply for the *z*-file. The *z*-file is created mode 444.

Limitations of the `get` Command

If the effective user has write permission (either explicitly or implicitly) in the directory containing the SCCS files, but the real user doesn't, only one file may be named when the `-e` option is used.

A.10. `help` — Ask for SCCS Help

`help` finds information to explain a message from a command or explain the use of a command. Zero or more arguments may be supplied. If no arguments are given, `help` prompts for one.

```
help [args]
```

The arguments may be either message numbers (which normally appear in parentheses following messages) or command names, of one of the following types:

type 1 Begins with non-numeric, ends in numerics. The non-numeric prefix is usually an abbreviation for the program or set of routines which produced the message (for example, `ge6`, for message 6 from the `get` command).

type 2 Does not contain numerics (as a command, such as `get`)

type 3 Is all numeric (for example, `212`)

The response of the program is the explanatory information related to the argument, if there is any.

When all else fails, try `help stuck`.

Example of `help`

The following asks for help on the `ge5` error message and information about the `rmdel` command:

```
tutorial% help ge5 rmdel
ge5:
"nonexistent sid"
The specified sid does not exist in the
given file.
Check for typos.

rmdel:
  rmdel -rSID name ...
tutorial%
```

Files Used

`/usr/lib/help`
directory containing files of message text.

A.11. `prs` — Print SCCS File

`prs` prints, on the standard output, parts or all of an SCCS file (see *sccsfile* (5)) in a user supplied format. If a directory is named, `prs` behaves as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the path name does not begin with `s.`), and unreadable files are silently ignored. If a name of `-` is given, the standard input is read, in which case each line is taken to be the name of an SCCS file or directory to be processed; non-SCCS files and unreadable files are silently ignored.

```
prs [ -d[dataspec] ] [ -r[SID] ] [ -e ] [ -l ] [ -a ] filename ...
```

prs Options

Options apply independently to each named file.

Output data specification

-d [*dataspec*]

Specifies the output data specification. The *dataspec* is a string consisting of SCCS file *data keywords* (see A.11.2) interspersed with optional user supplied text.

ID string

-r [*SID*]

Specifies the SCCS *ID*entification (ID) string of a delta for which information is desired. If no *SID* is specified, the *SID* of the most recently created delta is assumed.

Information on earlier deltas

-e Requests information for all deltas created *earlier* than and including the delta designated via the **-r** option.

Information on later deltas

-l Requests information for all deltas created *later* than and including the delta designated via the **-r** option.

Information for all deltas

-a Requests printing of information for both removed, that is, delta type = *R*, (see `rmDEL`) and existing, that is, delta type = *D*, deltas. If the **-a** option is not specified, information for existing deltas only is provided.

In the absence of the **-d** options, `prs` displays a default set of information consisting of: delta-type, release number and level number, date and time last changed, user-name of the person who changed the file, lines inserted, changed, and unchanged, the MR numbers, and the comments.

Data Keywords

Data keywords specify which parts of an SCCS file are to be retrieved and output. All parts of an SCCS file (see `scCSfile(5)`) have an associated data keyword. There is no limit on the number of times a data keyword may appear in a *dataspec*.

The information printed by `prs` consists of: 1) the user supplied text; and 2) appropriate values (extracted from the SCCS file) substituted for the recognized data keywords in the order of appearance in the *dataspec*. The format of a data keyword value is either *Simple* (*S*), in which keyword substitution is direct, or *Multi-line* (*M*), in which keyword substitution is followed by a carriage return.

User supplied text is any text other than recognized data keywords. A tab is specified by `\t` and carriage return/new-line is specified by `\n`.

Table A-3 SCCS Files Data Keywords

<i>Keyword</i>	<i>Data Item</i>	<i>File Section</i>	<i>Value</i>	<i>Format</i>
:Dt :	Delta information	Delta Table	See below*	S
:DL :	Delta line statistics	"	:Li:/:Ld:/:Lu:	S
:Li :	Lines inserted by Delta	"	nnnnn	S
:Ld :	Lines deleted by Delta	"	nnnnn	S
:Lu :	Lines unchanged by Delta	"	nnnnn	S
:DT :	Delta type	"	D or R	S
:I :	SCCS ID string (SID)	"	:R::L::B::S:	S
:R :	Release number	"	nnnn	S
:L :	Level number	"	nnnn	S
:B :	Branch number	"	nnnn	S
:S :	Sequence number	"	nnnn	S
:D :	Date Delta created	"	:Dy:/:Dm:/:Dd:	S
:Dy :	Year Delta created	"	nn	S
:Dm :	Month Delta created	"	nn	S
:Dd :	Day Delta created	"	nn	S
:T :	Time Delta created	"	:Th:::Tm:::Ts:	S
:Th :	Hour Delta created	"	nn	S
:Tm :	Minutes Delta created	"	nn	S
:Ts :	Seconds Delta created	"	nn	S
:P :	Programmer who created Delta	"	logname	S
:DS :	Delta sequence number	"	nnnn	S
:DP :	Predecessor Delta seq-no.	"	nnnn	S
:DI :	Seq-no. of deltas incl., excl., ignored	"	:Dn:/:Dx:/:Dg:	S
:Dn :	Deltas included (seq #)	"	:DS: :DS: ...	S
:Dx :	Deltas excluded (seq #)	"	:DS: :DS: ...	S
:Dg :	Deltas ignored (seq #)	"	:DS: :DS: ...	S
:MR :	MR numbers for delta	"	text	M
:C :	Comments for delta	"	text	M
:UN :	User names	User Names	text	M
:FL :	Flag list	Flags	text	M
:Y :	Module type flag	"	text	S
:MF :	MR validation flag	"	yes or no	S
:MP :	MR validation pgm name	"	text	S
:KF :	Keyword error/warning flag	"	yes or no	S
:BF :	Branch flag	"	yes or no	S
:J :	Joint edit flag	"	yes or no	S
:LK :	Locked releases	"	:R: ...	S
:Q :	User defined keyword	"	text	S
:M :	Module name	"	text	S
:FB :	Floor boundary	"	:R:	S
:CB :	Ceiling boundary	"	:R:	S
:Ds :	Default SID	"	:i:	S
:ND :	Null delta flag	"	yes or no	S
:FD :	File descriptive text	Comments	text	M
:BD :	Body	Body	text	M
:GB :	Gotten body	"	text	M
:W :	A form of <i>what</i> (1) string	N/A	:Z::M:~t:I:	S
:A :	A form of <i>what</i> (1) string	N/A	:Z::Y: :M: :I::Z:	S

Table A-3 *SCCS Files Data Keywords—Continued*

<i>Keyword</i>	<i>Data Item</i>	<i>File Section</i>	<i>Value</i>	<i>Format</i>
:Z:	<i>what(1)</i> string delimiter	N/A	@(#)	S
:F:	SCCS file name	N/A	text	S
:PN:	SCCS file path name	N/A	text	S

* :Dt: = :DT: :I: :D: :T: :P: :DS: :DP:

Examples of Using `prs`

```
tutorial% prs -d"Users and/or user IDs for :F: are:\n:UN:" s.file
```

may produce on the standard output:

```
Users and/or user IDs for s.file are:
xyz
131
abc
```

```
tutorial% prs -d"Newest delta for pgm :M:: :I: Created :D: By :P:" -r s.file
```

may produce on the standard output:

```
Newest delta for pgm main.c: 3.7 Created 77/12/1 By cas
```

As a *special case*:

```
tutorial% prs s.file
```

may produce on the standard output:

```
D 1.1 77/12/1 00:00:00 cas 1 000000/00000/00000
MRs:
bl78-12345
bl79-54321
COMMENTS:
this is the comment line for s.file initial delta
```

for each delta table entry of the “D” type. The only option argument allowed to be used with the *special case* is the `-a` option.

Files Used

`/tmp/pr?????`

A.12. `rmDEL` — Remove Delta from SCCS File

`rmDEL` removes the delta specified by the *SID* from each named SCCS *file*. The delta to be removed must be the newest (most recent) delta in its branch in the delta chain of each named SCCS file. In addition, the *SID* specified must *not* be that of a version being edited for the purpose of making a delta (that is, if a *p-file* (see `get`) exists for the named SCCS file, the *SID* specified must *not* appear in any entry of the *p-file*).

```
rmDEL -rSID filename . . .
```

If a directory is named, `rmDEL` behaves as though each file in the directory were specified as a named file, except that non-SCCS files (last component of the path name does not begin with `s.`) and unreadable files are silently ignored. If a name of `-` is given, the standard input is read; each line of the standard input is

taken to be the name of an SCCS file to be processed; non-SCCS files and unreadable files are silently ignored.

The exact permissions necessary to remove a delta are documented in the *Source Code Control System User's Guide*. Simply stated, they are either 1) if you make a delta you can remove it; or 2) if you own the file and directory you can remove a delta.

The delta to be removed must be a 'leaf' delta; that is, it must be the latest (most recently created) delta on its branch or on the trunk of the SCCS file tree. In Figure A-3, only deltas 1.3.1.2, 1.3.2.2, and 2.2 can be removed; once they are removed, deltas 1.3.2.1 and 2.1 can be removed, and so on.

To remove a delta, the effective user must have write permission in the directory containing the SCCS file. In addition, the real user must either have created the delta being removed, or be the owner of the SCCS file and its directory.

You must specify the *complete* SID of the delta to be removed, preceded by `-r`. The SID must have two components for a trunk delta, and four components for a branch delta. Thus:

```
tutorial% rmdel -r2.3 s.abc
```

removes (trunk) delta '2.3' of the SCCS file.

Before removing the delta, `rmdel` checks the following things:

1. the *release* number (R) of the given SID satisfies the relation:

$$\text{floor} \leq R \leq \text{ceiling}$$
2. the SID specified is *not* that of a version for which a `get` for editing has been executed and whose associated `delta` has not yet been made.
3. the login name or group ID of the user either appears in the file's *user list* or the *user list* is empty.
4. the release specified cannot be *locked* against editing (that is, if the `l` flag is set (see `admin`), the release specified *must* not be contained in the list).

If these conditions are satisfied, the delta is removed. Otherwise, processing is terminated.

After the specified delta has been removed, its type indicator in the *delta table* of the SCCS file is changed from 'D' (delta) to 'R' (removed).

Files Used

x-file (see `delta`)

z-file (see `delta`)

A.13. `sact` — Display SCCS Editing Activity

`sact` informs the user of any SCCS files which have had one or more `get -e` commands applied to them, that is, there are files out for editing, and deltas are pending. If a directory is named on the command line, `sact` behaves as though each file in the directory were specified as a named file, except that non-SCCS files and unreadable files are silently ignored. If a name of `-` is given, the standard input is read with each line being taken as the name of an SCCS file to be processed.

```
sact filename . . .
```

The output for each named file consists of five fields separated by spaces.

<i>Field Number</i>	<i>Meaning</i>
1	specifies the SID of a delta that currently exists in the SCCS file to which changes will be made to make the new delta.
2	specifies the SID for the new delta to be created.
3	contains the logname of the user who will make the delta (that is, executed a <code>get</code> for editing).
4	contains the date that <code>get -e</code> was executed.
5	contains the time that <code>get -e</code> was executed.

A.14. `sccsdiff` — Display Differences in SCCS Versions

`sccsdiff` compares two versions of an SCCS file and generates the differences between the two versions. Any number of SCCS files may be specified, but options apply to all files.

```
sccsdiff -rSID1 -rSID2 [ -p ] [ -sn ] filename . . .
```

`sccsdiff` Options

`-rSID?`

SID1 and *SID2* specify the deltas of an SCCS file that are to be compared. Versions are passed to `diff` in the order given.

`-p` pipe output for each file through `pr`.

`-sn`

n is the file segment size that `diff` will use. This is useful when the system load is high.

Files Used

`/tmp/get?????`

Temporary files

Diagnostics from `sccsdiff`

`file`: No differences

If the two versions are the same.

A.15. `unget` — Undo a Previous SCCS `get`

`Unget` undoes the effect of a `get -e` done prior to creating the intended new delta. If a directory is named, `unget` behaves as though each file in the directory were specified as a named file, except that non-SCCS files and unreadable files are silently ignored. If a name of `-` is given, the standard input is read with each line being taken as the name of an SCCS file to be processed.

```
unget [ -rSID ] [ -s ] [ -n ] filename . . .
```

`unget` Options

Options apply independently to each named file.

Delta to be removed

`-r SID`

Uniquely identifies which delta is no longer intended. (This would have been specified by `get` as the “new delta”). The `-r` option is necessary only if two or more outstanding `get`s for editing on the same SCCS file were done by the same person (login name). A diagnostic results if the specified `SID` is ambiguous, or if it is necessary but omitted from the command line.

Suppress delta ID

`-s` Suppress displaying the intended delta’s `SID`.

Retain gotten file

`-n` Retain the gotten file — it is normally removed from the current directory.

A.16. `val` — Validate SCCS File

`val` determines if the specified *file* is an SCCS file meeting the characteristics specified by the optional argument list. Arguments to `val` may appear in any order.

```
val -
    or
val [ -s ] [ -rSID ] [ -mname ] [ -ytype ] filename . . .
```

`val` has a special argument, `-`, which means read the standard input until an end-of-file condition is detected. Each line read is independently processed as if it were a command line argument list.

`val` generates diagnostic messages on the standard output for each command line and file processed and also returns a single 8-bit code upon exit as described below.

`val` Options

Options apply independently to each named file on the command line.

Suppress error messages

`-s` Silence diagnostic messages normally generated for errors detected while processing the specified files.

Delta number

`-r SID`

The argument value ID (SCCS IDentification String) is an SCCS delta number. A check is made to determine if the `SID` is ambiguous (for instance, `-r1` is ambiguous because it physically does not exist but implies 1.1, 1.2,

etc. which may exist) or invalid (for instance, `-r1.0` or `-r1.1.0` are invalid because neither case can exist as a valid delta number). If the SID is valid and not ambiguous, a check is made to determine if it actually exists.

Compare module names

`-m name`
name is compared with the SCCS `%M%` keyword in *file*.

Compare module types

`-y type`
type is compared with the SCCS `%Y%` keyword in *file*.

The 8-bit code returned by `val` is a disjunction of the possible errors, that is, can be interpreted as a bit string where (moving from left to right) set bits are interpreted as follows:

Table A-4 *Codes Returned from val Command*

<i>Bit</i>	<i>Meaning</i>
0	missing file argument
1	unknown or duplicate option
2	corrupted SCCS file
3	can't open file or file not SCCS
4	SID is invalid or ambiguous
5	SID does not exist
6	<code>%Y%</code> , <code>-y</code> mismatch
7	<code>%M%</code> , <code>-m</code> mismatch

Note that `val` can process two or more files on a given command line and in turn can process multiple command lines (when reading the standard input). In these cases an aggregate code is returned — logical OR of the codes generated for each command line and file processed.

Limitations of the `val` Command

`val` can process up to 50 files on a single command line. Any number above 50 produces a memory dump.

what — Identify SCCS Files

`what` finds SCCS identifying information within *any* specified UNIX file. `what` does not use any options, nor does it treat directory names and a name of `'-'` (a lone minus sign) in any special way, as do other SCCS commands.

`what` searches the given file(s) for all occurrences of the string `@(#)`, which is the replacement for the `%Z%` ID keyword (see `get`). `what` then displays whatever follows that string until the first double quote (`"`), greater than (`>`), backslash (`\`), newline, or (non-printing) NUL character.

As an example, let's begin with the SCCS file `s.program.c` (a C program), which contains the following line:

```
char id[] "%Z%M%:%I%";
```

We then do the following `get`:

```
tutorial% get -r3.4 s.prog.c
```

and finally compile the resulting *g-file* to produce `prog.o` and `a.out`. Using what as follows then displays:

```
tutorial% what prog.c prog.o a.out
prog.c:
  prog.c:3.4
prog.o:
  prog.c:3.4
a.out:
  prog.c:3.4
tutorial%
```

The string `what` searches for need not be inserted via an ID keyword of `get` — it may be inserted in any convenient manner.

A.17. SCCS Files

This section discusses several topics that must be considered before extensive use is made of SCCS. These topics deal with the protection mechanisms relied upon by SCCS, the format of SCCS files, and the recommended procedures for auditing SCCS files.

Protection

SCCS relies on the capabilities of the UNIX operating system for most of the protection mechanisms required to prevent unauthorized changes to SCCS files (that is, changes made by non-SCCS commands). The only protection features provided directly by SCCS are the *release lock* flag, the *release floor* and *ceiling* flags, and the *user list*.

New SCCS files created by `admin` are given mode 444 (read-only). It is best *not* to change this mode, as it prevents any direct modification of the files by non-SCCS commands.

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.

SCCS files must have only *one* link (name). Commands that modify SCCS files do so by creating a temporary copy of the file (called the *x-file*), and, upon completion of processing, remove the old file and rename the *x-file*. If the old file has more than one link, removing it and renaming the *x-file* would break the link. Rather than process such files, SCCS commands produce an error message. All SCCS files *must* have names that begin with 's.'

When only one user uses SCCS, the real and effective user IDs are the same, and that user ID owns the directories containing SCCS files. Therefore, SCCS may be used directly without any preliminary preparation.

However, in those situations in which several users with different user IDs are assigned responsibility for one SCCS file (for example, in large software

development projects), one user (equivalently, one user ID) must be chosen as the ‘owner’ of the SCCS files and as the one who will ‘administer’ them (for example, by using `admin`). This user is termed the *SCCS administrator* for that project. Because other users of SCCS do not have the same privileges and permissions as the SCCS administrator, they are not able to execute directly those commands that require write permission in the directory containing the SCCS files. Therefore, a project-dependent program is required to provide an interface to the `get`, `delta`, and, if desired, `rmDEL` and `cdc` commands.

The interface program must be owned by the SCCS administrator, and must have the *set-user-ID on execution* bit on (see `chmod(1)`), so that the effective user ID is the administrator’s user ID. This program’s function is to invoke the desired SCCS command and to cause it to *inherit* the privileges of the interface program for the duration of that command’s execution. In this manner, the owner of an SCCS file can modify it at will. Other users whose *login* names or *group* IDs are in the *user list* for that file (but who are *not* its owners) are given the necessary permissions only for the duration of the execution of the interface program, and are thus able to modify the SCCS files only through the use of `delta` and, possibly, `rmDEL` and `cdc`. The project-dependent interface program, as its name implies, must be custom-built for each project.

Layout of an SCCS File

SCCS files are composed of lines of ASCII text arranged in six parts, as follows:

Checksum	A line containing the ‘logical’ sum of all the characters of the file (<i>not</i> including this checksum itself).
Delta Table	Information about each delta, such as its type, SID, date and time of creation, and commentary included.
User Names	List of login names and/or group IDs of users who are allowed to modify the file by adding or removing deltas.
Flags	Indicators that control certain actions of various SCCS commands.
Descriptive Text	Text provided by the user; usually a summary of the contents and purpose of the file.
Body	Actual text that is being administered by SCCS, intermixed with internal SCCS control lines.

Detailed information about the contents of the various sections of the file may be found in `sccsfile(5)`. In the following, the *checksum* is the only portion of the file discussed.

Because SCCS files are ASCII files, they may be processed by various UNIX commands: editors such as `vi(1)`, text processing programs such as `grep(1)`, `awk(1)`, and `cat(1)`, and so on. This is quite useful when an SCCS file must be modified manually (for example, when the time and date of a delta was recorded incorrectly because the system clock was set incorrectly), or when one wants to simply ‘look’ at the file.

CAUTION Extreme care should be exercised when modifying SCCS files with non-SCCS commands.

Auditing

On rare occasions, perhaps due to an operating system or hardware malfunction, all or part of an SCCS file is destroyed. SCCS commands (like most UNIX commands) display an error message when a file does not exist. In addition, SCCS commands use the *checksum* stored in the SCCS file to determine whether a file has been *corrupted* since it was last accessed (has lost data, or has been changed). The *only* SCCS command which will process a corrupted SCCS file is `admin` with the `-h` or `-z` options. This is discussed below.

SCCS files should be audited (checked) for possible corruptions on a regular basis. The simplest and fastest way to audit such files is to use `admin` with the `-h` option on them:

```
tutorial% admin -h s.file1 s.file2 ...
           or
tutorial% admin -h directory1 directory2 ...
```

If the new checksum of any file is not equal to the checksum in the first line of that file, the message

```
corrupted file (co6)
```

is produced for that file. This process continues until all files have been examined. When examining directories (as in the second example above), the process just described does not detect *missing* files. A simple way to detect whether any files are missing from a directory is to periodically list the contents of the directory (using `ls(1)`), and compare the current listing with the previous one. Any file which appears on the previous list but not the current one has been removed by some means.

When a file has been corrupted, the appropriate method of restoration depends upon the extent of the corruption. If damage is extensive, the best solution is to restore the file from a backup copy. When damage is minor, repairing the file with your favorite text editor may be possible. If you do repair the file with the system's text processing capabilities, you must use `admin` with the `-z` option to recompute the checksum to bring it into agreement with the actual contents of the file:

```
tutorial% admin -z s.file
```

After this command is executed on a file, any corruption which may have existed in that file will no longer be detectable.

B

Bibliography and Credits

Bibliography and Credits	231
--------------------------------	-----

Bibliography and Credits

This appendix contains acknowledgements from the original authors of the papers in this manual, plus references to other possibly interesting literature for the diligent scholar.

Reference numbers here are keyed to the reference numbers in the specific chapters.

UNIX Programming References

- [1] Sun Microsystems Reference Manuals: *Commands Reference Manual* and *System Interface Manual*.
- [2] B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Prentice-Hall, Inc., 1978.
- [3] B. W. Kernighan, *UNIX for Beginners — Second Edition*, Bell Laboratories, 1978. Reprinted in the *Sun Tutorial for Beginners Manual*.

make Acknowledgments and References

I would like to thank S. C. Johnson for suggesting this approach to program maintenance control. I would like to thank S. C. Johnson and H. Gajewska for being the prime guinea pigs during development of *make*.

1. S. C. Johnson, 'Yacc — Yet Another Compiler-Compiler', Bell Laboratories Computing Science Technical Report #32, July 1978.
2. M. E. Lesk, 'Lex — A Lexical Analyzer Generator', Computing Science Technical Report #39, October 1975.

Performance Analysis Acknowledgments and References

- [1] *gprof: a Call Graph Execution Profiler*, by Susan L. Graham, Peter B. Kessler, and Marshall Kirk McKusick, Computer Science Division, Electrical Engineering and Computer Science Department, University of California, Berkeley, California 94720. This paper is for the scholar interested in the theory behind call-graph profiling.

m4 Acknowledgments and References

We are indebted to Rick Becker, John Chambers, Doug McIlroy, and especially Jim Weythman, whose pioneering use of *m4* has led to several valuable improvements. We are also deeply grateful to Weythman for several substantial contributions to the code.

The *m4* macro processor is an extension of a macro processor called M3 which was written by D. M. Ritchie for the AP-3 minicomputer; M3 was in turn based on a macro processor implemented for [1]. Readers unfamiliar with the basic ideas of macro processing may wish to read some of the discussion there.

- [1] B. W. Kernighan and P. J. Plauger, *Software Tools* Addison-Wesley, Inc., 1976.

lex Acknowledgments and References

As should be obvious from the above, the outside of *lex* is patterned on *yacc* and the inside on Aho's string matching routines. Therefore, both S. C. Johnson and A. V. Aho are really originators of much of *lex*, as well as debuggers of it. Many thanks are due to both.

The code of the current version of *lex* was designed, written, and debugged by Eric Schmidt.

1. B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Prentice-Hall, N. J. (1978).
2. B. W. Kernighan, *Ratfor: A Preprocessor for a Rational FORTRAN*, Software — Practice and Experience, 5, pp. 395-496 (1975).
3. S. C. Johnson, *Yacc: Yet Another Compiler Compiler*, Computing Science Technical Report No. 32, 1975, Murray Hill.
4. A. V. Aho and M. J. Corasick, *Efficient String Matching: An Aid to Bibliographic Search*, Comm. ACM 18, 333-340 (1975).
5. See the papers on *ex* and *vi* in *Editing and Text Processing on the Sun Workstation*.

**lint Acknowledgments and
References**

6. M. E. Lesk, *The Portable C Library*, Computing Science Technical Report No. 31, Murray Hill.
1. B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Prentice-Hall, N. J. (1978).
2. S. C. Johnson, *Yacc: Yet Another Compiler-Compiler*, Comp. Sci. Tech. Rep. No. 32, Bell Laboratories, Murray Hill, New Jersey (July 1975).
3. M. E. Lesk, *Lex — A Lexical Analyzer Generator*, Comp. Sci. Tech. Rep. No. 39, Bell Laboratories, Murray Hill, New Jersey (October 1975).
4. S. C. Johnson and D. M. Ritchie, 'UNIX Time-Sharing System: Portability of C Programs and the UNIX System,' *Bell Sys. Tech. J.* **57**(6) pp. 2021-2048 (1978).
5. S. C. Johnson, 'A Portable Compiler: Theory and Practice,' *Proc. 5th ACM Symp. on Principles of Programming Languages*, (January 1978).

Index

A

acknowledgements, 231 *thru* 233
actions
 in `lex`, 126 *thru* 129
actions for `yacc`, 147
`admin` — administer SCCS, 190 *thru* 197
ambiguity in `yacc`, 155
ambiguous source rules
 in `lex`, 129 *thru* 131
`ARGSUSED` — lint control, 44
associativity in `yacc`
 `%left`, 159
 `%nonassoc`, 159
 `%right`, 159
auditing SCCS files, 228

B

backup files with SCCS commands
 `x-file`, 190
 `z-file`, 190
basic specifications for `yacc`, 145
bibliography, 231 *thru* 233
branch number in SCCS file, 187
build programs — `make`, 49 *thru* 68
built-in `m4` macros
 `changequote`, 110
 `define`, 108
 `divert`, 112
 `divnum`, 113
 `dnl`, 114
 `dumpdef`, 115
 `errprint`, 115
 `eval`, 111
 `ifdef`, 110
 `ifelse`, 113
 `include`, 112
 `incr`, 111
 `index`, 114
 `len`, 114
 `mktemp`, 113
 `sinclude`, 112
 `substr`, 114
 `syscmd`, 113
 `translit`, 114
 `undefine`, 110
 `undivert`, 113
built-in rules

built-in rules, *continued*
 double, 66
 single, 65
built-in suffix list, 65 *thru* 68
built-in transformation rules, 65 *thru* 68

C

C language tools, 33 *thru* 46
 `ctags` — `ctags` used, 33 *thru* 34
 `lint` — check C programs, 35 *thru* 46
call graph profile — `gprof`, 98 *thru* 100
`cdc` — change delta commentary, 197 *thru* 198
change
 delta commentary, 197 *thru* 198
`changequote` built-in `m4` macro, 110
character set
 in `lex`, 138
checksum in SCCS files, 227
code coverage — `tcov`, 100 *thru* 103
`comb` — combine deltas, 198 *thru* 199
combine SCCS deltas, 198 *thru* 199
command conventions for SCCS, 189 *thru* 190
commands in SCCS
 backup files with SCCS commands, 190
 file arguments, 189
 flags, 190
 options, 189
comments in `m4`, 110
comments in `make`, 52
compare versions of SCCS file — `sccsdiff`, 223
compiler generators
 `lex` lexical generator, 119 *thru* 140
compiler-compiler
 `yacc`, 143 *thru* 182
conflicts in `yacc`, 156
 disambiguating rules, 156
 precedence, 159
 reduce/reduce conflicts, 156
 shift/reduce conflicts, 156
continuation lines in `make`, 52
create
 delta, 200 *thru* 204
 SCCS data bases, 190 *thru* 197
 SCCS delta, 200 *thru* 204
create SCCS database file — `sccs create`, 75
credits, 231 *thru* 233

ctags — ctags used, 33 *thru* 34

D

data keywords in SCCS files, 218 *thru* 220
 define built-in m4 macro, 108
 defining make macros, 53 *thru* 54
 definitions
 of lex source, 131 *thru* 132
 delta, 185, 74
 change commentary, 197 *thru* 198
 combine, 198 *thru* 199
 remove — rmdel, 221 *thru* 222
 make SCCS delta — sccs delta, 76, 200 *thru* 204
 delta table in SCCS files, 227
 dependency lines in make, 54 *thru* 56
 descriptive text in SCCS files, 227
 disambiguating rules in yacc, 156
 display
 display SCCS file editing status — sact, 223
 SCCS history — prs, 217 *thru* 221
 divert built-in m4 macro, 112
 divnum built-in m4 macro, 113
 dn1 built-in m4 macro, 114
 double suffix rules, 66
 dumpdef built-in m4 macro, 115

E

get SCCS file for editing — sccs edit, 76
 entries in make, 54 *thru* 56
 errprint built-in m4 macro, 115
 eval built-in m4 macro, 111
 examples
 of lex, 133 *thru* 136
 of make, 61 *thru* 62

F

file arguments to SCCS commands, 189
 flags in SCCS files, 193 *thru* 194, 227
 flags to SCCS commands, 190

G

get — get SCCS file, 204 *thru* 216
 get SCCS file for compiling — sccs get, 75
 gprof — call graph, 98 *thru* 100

H

help — get SCCS help, 216 *thru* 217
 high-level SCCS interface, 71 *thru* 90

I

ID keywords, 74
 Id keywords
 inserting, 77
 placement of, 78
 identification keywords in SCCS files, 208 *thru* 209
 identify SCCS file — what, 225 *thru* 226
 ifdef built-in m4 macro, 110
 ifndef built-in m4 macro, 113
 implicit make macros

implicit make macros, *continued*

\$*, 56
 \$<, 56
 \$?, 56
 \$@, 56

implicit rules in make, 59 *thru* 61
 include built-in m4 macro, 112
 include files in make, 53
 incr built-in m4 macro, 111
 index built-in m4 macro, 114

L

language tools
 lint — check C programs, 35 *thru* 46
 yacc compiler-compiler, 143 *thru* 182
 layout of SCCS files, 227 *thru* 228
 body, 227
 checksum, 227
 delta table, 227
 descriptive text, 227
 flags, 227
 user names, 227
 left association in yacc, 155
 left context-sensitivity
 in lex, 136 *thru* 138
 len built-in m4 macro, 114
 level number in SCCS file, 185
 lex
 actions, 126 *thru* 129
 ambiguous source rules, 129 *thru* 131
 character set, 138
 examples, 133 *thru* 136
 left context-sensitivity, 136 *thru* 138
 regular expressions, 123 *thru* 126
 source definitions, 131 *thru* 132
 source format, 122
 source format summary, 138 *thru* 140
 usage, 132
 with yacc, 133
 lex lexical generator, 119 *thru* 140
 lex regular expressions
 arbitrary character, 124
 character classes, 124
 context sensitivity, 125
 operators, 123
 optional expressions, 124
 repeated expressions, 125
 repetitions and definition, 126
 lexical analysis for yacc, 149
 lint — C program checker, 35 *thru* 46
 lint controls, 44
 ARGSUSED, 44
 LINTLIBRARY, 44
 NOSTRICT, 44
 NOTREACHED, 44
 VARARGS, 44
 VARARGS2, 44
 lint library directive
 LINTLIBRARY, 45
 lint library directives, 45
 lint options, 46

LINTLIBRARY — lint control, 44
 LINTLIBRARY — lint library directives, 45
 low-level SCCS interface, 185 *thru* 228

M

m4 built-in macros
 changequote, 110
 define, 108
 divert, 112
 divnum, 113
 dnl, 114
 dumpdef, 115
 errprint, 115
 eval, 111
 ifdef, 110
 ifelse, 113
 include, 112
 incr, 111
 index, 114
 len, 114
 mktemp, 113
 sinclude, 112
 substr, 114
 syscmd, 113
 translit, 114
 undefine, 110
 undivert, 113

m4 macro processor, 107 *thru* 116

macros in make, 53 *thru* 54

maintain programs — make, 49 *thru* 68

make
 comment convention, 52
 continuation lines, 52
 delta, 200 *thru* 204
 dependency lines, 54 *thru* 56
 entries, 54 *thru* 56
 implicit rules, 59 *thru* 61
 include files, 53
 macros, 53 *thru* 54
 rules lines, 54 *thru* 56
 SCCS delta, 200 *thru* 204
 target lines, 54 *thru* 56
 using command, 56 *thru* 58

make — build programs, 49 *thru* 68

make implicit macros
 \$*, 56
 \$<, 56
 \$?, 56
 \$@, 56

make options, 58

mktemp built-in m4 macro, 113

N

NOSTRICT — lint control, 44
 NOTREACHED — lint control, 44

O

options
 lint, 46
 make, 58

options to SCCS commands, 189

P

parser generators
 yacc, 143 *thru* 182

performance analysis, 93 *thru* 103
 gprof — call graph, 98 *thru* 100
 prof — profile, 96 *thru* 98
 tcov — code coverage, 100 *thru* 103
 time — time used, 93 *thru* 96

precedence in yacc, 159
 %prec, 160

preparing yacc specifications, 164 *thru* 166

prof — profile, 96 *thru* 98

profile
 call graph — gprof, 98 *thru* 100

programming tools
 lint — check C programs, 35 *thru* 46
 make — build programs, 49 *thru* 68
 yacc compiler-compiler, 143 *thru* 182

protections on SCCS files, 226 *thru* 227

prs — display SCCS history, 217 *thru* 221

Q

quoting in m4, 109 *thru* 110

R

reduce/reduce conflicts in yacc, 156

references, 231 *thru* 233

regenerate programs — make, 49 *thru* 68

regular expressions
 in lex, 123 *thru* 126

release number in SCCS file, 185

remove delta from SCCS file — rmdel, 221 *thru* 222

right association in yacc, 155

rmdel — remove delta from SCCS file, 221 *thru* 222

rules, 65
 double, 66
 single, 65

rules lines in make, 54 *thru* 56

S

s-file, 74

SCCS backup files, 190
 x-file, 190
 z-file, 190

SCCS branch number, 187

SCCS command conventions, 189 *thru* 190

SCCS commands
 admin — administer SCCS, 190 *thru* 197
 cdc — change delta commentary, 197 *thru* 198
 comb — combine deltas, 198 *thru* 199
 file arguments, 189
 flags, 190
 get — get SCCS file, 204 *thru* 216
 help — get SCCS help, 216 *thru* 217
 options, 189
 prs — display SCCS history, 217 *thru* 221
 rmdel — remove delta, 221 *thru* 222
 sact — display SCCS file editing status, 223
 sccsdiff — compare versions of SCCS file, 223
 unget — unget SCCS file, 224
 val — validate SCCS file, 224 *thru* 225

SCCS commands, *continued*

what — identify SCCS file, 225 *thru* 226
sccs create — create SCCS database file, 75
 SCCS data keywords, 218 *thru* 220
sccs delta — make SCCS delta, 76
 SCCS delta
 change commentary, 197 *thru* 198
 combine, 198 *thru* 199
 create, 200 *thru* 204
 remove — *rmDEL*, 221 *thru* 222
sccs edit — get SCCS file for editing, 76
 SCCS files, 226 *thru* 228
 auditing, 228
 body, 227
 checksum, 227
 delta table, 227
 descriptive text, 227
 flags, 227
 flags in, 193 *thru* 194
 layout, 227 *thru* 228
 protection, 226 *thru* 227
 user names, 227
 SCCS front end — *sccs*, 71 *thru* 90
sccs get — get SCCS file for compiling, 75
 SCCS history
 display — *prs*, 217 *thru* 221
 SCCS ID keywords, 74
 SCCS identification keywords, 208 *thru* 209
 SCCS identification string, 185
sccs info — who's editing files, 77
 SCCS level number, 185
 SCCS release number, 185
 SCCS sequence number, 187
sccsdiff — compare versions of SCCS file, 223
 sequence number in SCCS file, 187
 shift/reduce conflicts in *yacc*, 156
 SID, 185, 74
sinclude built-in *m4* macro, 112
 single suffix rules, 65
 source code control
 high-level SCCS interface, 71 *thru* 90
 low-level SCCS interface, 185 *thru* 228
 source definitions
 in *lex*, 131 *thru* 132
 standard suffix list, 65 *thru* 68
 statement analysis — *tcov*, 100 *thru* 103
substr built-in *m4* macro, 114
 suffix rules
 double, 66
 single, 65
 suffixes, 65
 summary of *lex* source format, 138 *thru* 140
 summary of SCCS commands, 185 *thru* 228
syscmd built-in *m4* macro, 113

T

target lines in *make*, 54 *thru* 56
tcov — code coverage, 100 *thru* 103
time — time used, 93 *thru* 96
 transformation rules, 65

translit built-in *m4* macro, 114
 trunk of SCCS tree, 186

U

undefine built-in *m4* macro, 110
undivert built-in *m4* macro, 113
unget — *unget* SCCS file, 224
 update programs — *make*, 49 *thru* 68
 user names in SCCS files, 227
 user-level SCCS interface, 71 *thru* 90
 using *lex*, 132
 using *make*, 56 *thru* 58

V

val — validate SCCS file, 224 *thru* 225
 validate SCCS file — *val*, 224 *thru* 225
 VARARGS — lint control, 44
 VARARGS2 — lint control, 44

W

what — identify SCCS file, 225 *thru* 226
 who's editing SCCS files — *sccs info*, 77

X

x-file in SCCS commands, 190

Y

yacc
 actions, 147
 ambiguity, 155
 basic specifications, 145
 conflicts, 156
 disambiguating rules, 156
 left association, 155
 lexical analysis, 149
 precedence, 159
 reduce/reduce conflicts, 156
 right association, 155
 shift/reduce conflicts, 156
yacc associativity
 %left, 159
 %nonassoc, 159
 %right, 159
yacc compiler-compiler, 143 *thru* 182
yacc precedence
 %prec, 160
yacc specifications
 preparing, 164 *thru* 166

Z

z-file in SCCS commands, 190

Revision History

<i>Rev</i>	<i>Date</i>	<i>Comments</i>
A	15th July 1983	First release of this Manual.
B	15th August 1983	Second Release of this manual entailed a complete reorganization and some rewriting of the individual articles.
C	1st November 1983	Third Release of this manual entailed minor corrections and updates.
D	7th January 1984	Added chapter on Shell Programming. Added chapter on ADB. Many minor corrections and updates.
E	15 May 1985	Many minor corrections and updates. Extracted Assembly Language Reference Manual to make a separate manual.
F	15 February 1986	Previous placeholder for ADB gone — a complete debug tools manual was written for release 3.0. Chapters on programming the Bourne Shell and programming the C Shell now part of the manual entitled: <i>Doing More with UNIX — Beginner's Guide</i> . Chapters on <code>bc</code> and <code>dc</code> now part of the manual entitled: <i>Games, Demos, and Other Pursuits — Beginner's Guide</i> . Old chapter on <code>lint</code> folded into a chapter on tools for programming with the C language. Chapter on <code>make</code> upgraded to reflect use of System V augmented <code>make</code> . Added new chapters on UNIX programming tools by extracting relevant narrative material from out of the 'man' pages. Low-level SCCS documentation for the 'raw' SCCS interface was moved to an appendix. SCCS narrative material from the SCCS 'man' pages incorporated to form a reference summary. Acknowledgements and references placed in an appendix to get them out of the way. Created an index. Many minor corrections and updates.

Notes

Notes

Notes

Notes

Notes