

# Software Development System

---

---

Vol. II

# MICROPORT SYSTEMS

The material contained in this manual was reprinted with permission from AT&T and is comprised of excerpts from the following AT&T manuals.

<b>*UNIX System V - Release 2.0 Support Tools Guide</b>	April 1984 307-108, Issue 2
<b>UNIX System V - Release 2.0 Supplement</b> †INTEL Processors	September 1985 307-624, Issue 2
<b>UNIX System V - Release 2.0 Programming Guide</b>	April 1984 307-103, Issue 2
<b>UNIX System V - Release 2.0 Programmer Reference Manual</b> INTEL Processors	March 1985 307-627, Issue 1

\*UNIX is a trademark of AT&T Bell Laboratories

†INTEL is a trademark of Intel Corporation

Copyright © 1984, 1985 by AT&T

All rights reserved

Printed in U.S.A.

DIABLO is a registered trademark of Xerox Corporation

UNIX is a trademark of AT&T Bell Laboratories

iAPX 286 is a trademark of Intel Corporation

DOCUMENTER'S WORKBENCH is a trademark of AT&T

PDP and VAX are trademarks of Digital Equipment Corporation

HP is a trademark of Hewlett-Packard, Inc.

TEKTRONIX is a registered trademark of Tektronix, Inc.

TELETYPE is a trademark of AT&T Teletype Corporation

VERSATEC is a registered trademark of Versatec Corporation

# QUICK REFERENCE GUIDE

C LIBRARIES	C-17
OBJECT & MATH LIBRARIES	C-18
LEX ANALYZER GEN (LEX)	C-19
YACC	C-20
SYSTEM CALLS	S-2
SUBROUTINES & LIBRARIES	S-3
MATH LIBRARY FUNCTIONS	S-3M
MISC. LIBRARY FUNCTIONS	S-3X
FORTTRAN LIBRARY FUNCTIONS	S-3F
FILE FORMATS	S-4
MISCELLANY	S-5

# Table of Contents

---

---

TITLE	CHAPTER
C LIBRARIES.....	17
OBJECT AND MATH LIBRARIES.....	18
LEX ANALYZER GEN (LEX).....	19
YACC.....	20

TITLE	SECTION
SYSTEM CALLS.....	2
SUBROUTINES AND LIBRARIES.....	3
MATH LIBRARY FUNCTIONS.....	3M
MISC. LIBRARY FUNCTIONS.....	3X
FORTRAN LIBRARY FUNCTIONS.....	3F
FILE FORMATS.....	4
MISCELLANY.....	5

## Chapter 17

# C LIBRARIES

### GENERAL

This chapter and the next describe the libraries that are supported on the UNIX operating system. A library is a collection of related functions and/or declarations that simplify programming effort by linking only what is needed, allowing use of locally produced functions, etc. All of the functions described are also described in section 3, 3M, 3X and 3F of the Software Development System manual. Most of the declarations described are in section 5. The three main libraries on the UNIX system are:

- C library** This is the basic library for C language programs. The C library is composed of functions and declarations used for file access, string testing and manipulation, character testing and manipulation, memory allocation, and other functions. This library is described later in this chapter.
- Object file** This library provides functions for the access and manipulation of object files. This library is described in the next chapter.
- Math library** This library provides exponential, bessel functions, logarithmic, hyperbolic, and trigonometric functions. This library is described in the next chapter.

Some libraries consist of two portions - functions and declarations. In some cases, the user must request that the functions (and/or declarations) of a specific library be included in a program being compiled. In other cases, the functions (and/or declarations) are included automatically.

### **Including Functions**

When a program is being compiled, the compiler will automatically search the C language library to locate and include functions that are used in the program. This is the case only for the C library and no other library. In order for the compiler to locate and include functions from other libraries, the user must specify these libraries on the command line for the compiler. For example, when using functions of the math library, the user must request that the math library be searched by including the argument **-lm** on the command line, such as:

```
cc file.c -lm
```

The argument **-lm** must come after all files that reference functions in the math library in order for the link editor to know which functions to include in the a.out file.

This method should be used for all functions that are not part of the C language library.

### **Including Declarations**

Some functions require a set of declarations in order to operate properly. A set of declarations is stored in a file under the */usr/include* directory. These files are referred to as *header files*. In order to include a certain header file, the user must specify this request within the C language program. The request is in the form:

```
#include <file.h>
```

where *file.h* is the name of the file. Since the header files define the type of the functions and various preprocessor constants, they must be included before invoking the functions they declare.

The remainder of this chapter describes the functions and header files of the C Library. The description of the library begins with the actions required by the user to include the functions and/or header files in a program being compiled (if any). Following the description

of the actions required is information in three-column format of the form:

<b>function</b>	<b>reference(N)</b>	<b>Brief description.</b>
-----------------	---------------------	---------------------------

The functions are grouped by type while the reference refers to section 'N' in the Software Development System manual. Following this, are descriptions of the header files associated with these functions (if any).

## THE C LIBRARY

~~The C library consists of several types of functions. All the functions of the C library are loaded automatically by the compiler. Various declarations must sometimes be included by the user as required. The functions of the C library are divided into the following types:~~

- Input/output control
- String manipulation
- Character manipulation
- Time functions
- Miscellaneous functions.

### Input/Output Control

These functions of the C library are automatically included as needed during the compiling of a C language program. No command line request is needed.

The header file required by the input/output functions should be included in the program being compiled. This is accomplished by including the line:

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

near the beginning of each file that references an input or output function.

The input/output functions are grouped into the following categories:

- File access
- File status
- Input
- Output
- Miscellaneous.

### File Access Functions

<i><b>FUNCTION</b></i>	<i><b>REFERENCE</b></i>	<i><b>BRIEF DESCRIPTION</b></i>
<b>fclose</b>	<b>fclose(3S)</b>	Close an open stream.
<b>fdopen</b>	<b>fopen(3S)</b>	Associate stream with an <b>open(2)</b> ed file.
<b>fileno</b>	<b>ferror(3S)</b>	File descriptor associated with an open stream.
<b>fopen</b>	<b>fopen(3S)</b>	Open a file with specified permissions. <b>Fopen</b> returns a pointer to a stream which is used in subsequent references to the file.
<b>freopen</b>	<b>fopen(3S)</b>	Substitute named file in place of open stream.
<b>fseek</b>	<b>fseek(3S)</b>	Reposition the file pointer.
<b>pclose</b>	<b>popen(3S)</b>	Close a stream opened by <b>popen</b> .
<b>popen</b>	<b>popen(3S)</b>	Create pipe as a stream between calling process and command.



<b>rewind</b>	<b>fseek(3S)</b>	Reposition file pointer at beginning of file.
<b>setbuf</b>	<b>setbuf(3S)</b>	Assign buffering to stream.
<b>vsetbuf</b>	<b>setbuf(3S)</b>	Similar to <b>setbuf</b> , but allowing finer control.

### File Status Functions

<b><i>FUNCTION</i></b>	<b><i>REFERENCE</i></b>	<b><i>BRIEF DESCRIPTION</i></b>
<b>clearerr</b>	<b>ferror(3S)</b>	Reset error condition on stream.
<b>feof</b>	<b>ferror(3S)</b>	Test for "end of file" on stream.
<b>ferror</b>	<b>ferror(3S)</b>	Test for error condition on stream.
<b>ftell</b>	<b>fseek(3S)</b>	Return current position in the file.

### Input Functions

<b><i>FUNCTION</i></b>	<b><i>REFERENCE</i></b>	<b><i>BRIEF DESCRIPTION</i></b>
<b>fgetc</b>	<b>getc(3S)</b>	True function for <b>getc(3S)</b> .
<b>fgets</b>	<b>gets(3S)</b>	Read string from stream.
<b>fread</b>	<b>fread(3S)</b>	General buffered read from stream.

## C LIBRARIES

<b>fscanf</b>	<b>scanf(3S)</b>	Formatted read from stream.
<b>getc</b>	<b>getc(3S)</b>	Read character from stream.
<b>getchar</b>	<b>getc(3S)</b>	Read character from standard input.
<b>gets</b>	<b>gets(3S)</b>	Read string from standard input.
<b>getw</b>	<b>getc(3S)</b>	Read word from stream.
<b>scanf</b>	<b>scanf(3S)</b>	Read using format from standard input.
<b>sscanf</b>	<b>scanf(3S)</b>	Formatted from string.
<b>ungetc</b>	<b>ungetc(3S)</b>	Put back one character on stream.

### Output Functions

<b><i>FUNCTION</i></b>	<b><i>REFERENCE</i></b>	<b><i>BRIEF DESCRIPTION</i></b>
<b>fflush</b>	<b>fclose(3S)</b>	Write all currently buffered characters from stream.
<b>fprintf</b>	<b>printf(3S)</b>	Formatted write to stream.
<b>fputc</b>	<b>putc(3S)</b>	True function for <b>putc(3S)</b> .
<b>fputs</b>	<b>puts(3S)</b>	Write string to stream.
<b>fwrite</b>	<b>fread(3S)</b>	General buffered write to stream.

<b>printf</b>	<b>printf(3S)</b>	Print using format to standard output.
<b>putc</b>	<b>putc(3S)</b>	Write character to standard output.
<b>putchar</b>	<b>putc(3S)</b>	Write character to standard output.
<b>puts</b>	<b>puts(3S)</b>	Write string to standard output.
<b>putw</b>	<b>putc(3S)</b>	Write word to stream.
<b>sprintf</b>	<b>printf(3S)</b>	Formatted write to string.
<b>vfprintf</b>	<b>vprintf(3C)</b>	Print using format to stream by <b>varargs(5)</b> argument list.
<b>vprintf</b>	<b>vprintf(3C)</b>	Print using format to standard output by <b>varargs(5)</b> argument list.
<b>vsprintf</b>	<b>vprintf(3C)</b>	Print using format to stream string by <b>varargs(5)</b> argument list.

### Miscellaneous Functions

<b>FUNCTION</b>	<b>REFERENCE</b>	<b>BRIEF DESCRIPTION</b>
<b>ctermid</b>	<b>ctermid(3S)</b>	Return file name for controlling terminal.
<b>cuserid</b>	<b>cuserid(3S)</b>	Return login name for owner of current process.
<b>system</b>	<b>system(3S)</b>	Execute shell command.

## C LIBRARIES

<b>tempnam</b>	<b>tempnam</b> (3S)	Create temporary file name using directory and prefix.
<b>tmpnam</b>	<b>tmpnam</b> (3S)	Create temporary file name.
<b>tmpfile</b>	<b>tmpfile</b> (3S)	Create temporary file.

### String Manipulation Functions

These functions are used to locate characters within a string, copy, concatenate, and compare strings. These functions are automatically located and loaded during the compiling of a C language program. No command line request is needed since these functions are part of the C library. The string manipulation functions are declared in a header file that may be included in the program being compiled. This is accomplished by including the line:

```
#include <string.h>
```

near the beginning of each file that uses one of these functions.

<b><i>FUNCTION</i></b>	<b><i>REFERENCE</i></b>	<b><i>BRIEF DESCRIPTION</i></b>
<b>strcat</b>	<b>string</b> (3C)	Concatenate two strings.
<b>strchr</b>	<b>string</b> (3C)	Search string for character.
<b>strcmp</b>	<b>string</b> (3C)	Compares two strings.
<b>strcpy</b>	<b>string</b> (3C)	Copy string.
<b>strcspn</b>	<b>string</b> (3C)	Length of initial string not containing set of characters.
<b>strlen</b>	<b>string</b> (3C)	Length of string.

<b>strncat</b>	<b>string(3C)</b>	Concatenate two strings with a maximum length.
<b>strncmp</b>	<b>string(3C)</b>	Compares two strings with a maximum length.
<b>strncpy</b>	<b>string(3C)</b>	Copy string over string with a maximum length.
<b>strpbrk</b>	<b>string(3C)</b>	Search string for any set of characters.
<b>strrchr</b>	<b>string(3C)</b>	Search string backwards for character.
<b>strspn</b>	<b>string(3C)</b>	Length of initial string containing set of characters.
<b>strtok</b>	<b>string(3C)</b>	Search string for token separated by any of a set of characters.

### Character Manipulation

The following functions and declarations are used for testing and translating ASCII characters. These functions are located and loaded automatically during the compiling of a C language program. No command line request is needed since these functions are part of the C library.

The declarations associated with these functions should be included in the program being compiled. This is accomplished by including the line:

```
#include <ctype.h>
```

near the beginning of the file being compiled.

**Character Testing Functions**

These functions can be used to identify characters as uppercase or lowercase letters, digits, punctuation, etc.

<b><i>FUNCTION</i></b>	<b><i>REFERENCE</i></b>	<b><i>BRIEF DESCRIPTION</i></b>
<b>isalnum</b>	<b>ctype(3C)</b>	Is character alphanumeric?
<b>isalpha</b>	<b>ctype(3C)</b>	Is character alphabetic?
<b>isascii</b>	<b>ctype(3C)</b>	Is integer ASCII character?
<b>isctrl</b>	<b>ctype(3C)</b>	Is character a control character?
<b>isdigit</b>	<b>ctype(3C)</b>	Is character a digit?
<b>isgraph</b>	<b>ctype(3C)</b>	Is character a printable character?
<b>islower</b>	<b>ctype(3C)</b>	Is character a lowercase letter?
<b>isprint</b>	<b>ctype(3C)</b>	Is character a printing character including space?
<b>ispunct</b>	<b>ctype(3C)</b>	Is character a punctuation character?
<b>isspace</b>	<b>ctype(3C)</b>	Is character a white space character?
<b>isupper</b>	<b>ctype(3C)</b>	Is character an uppercase letter?
<b>isxdigit</b>	<b>ctype(3C)</b>	Is character a hex digit?

### Character Translation Functions

These functions provide translation of uppercase to lowercase, lowercase to uppercase, and integer to ASCII.

<i>FUNCTION</i>	<i>REFERENCE</i>	<i>BRIEF DESCRIPTION</i>
<b>toascii</b>	<b>conv(3C)</b>	Convert integer to ASCII character.
<b>tolower</b>	<b>conv(3C)</b>	Convert character to lowercase.
<b>toupper</b>	<b>conv(3C)</b>	Convert character to uppercase.

### Time Functions

These functions are used for accessing and reformatting the systems idea of the current date and time. These functions are located and loaded automatically during the compiling of a C language program. No command line request is needed since these functions are part of the C library.

The header file associated with these functions should be included in the program being compiled. This is accomplished by including the line:

```
#include <time.h>
```

near the beginning of any file using the time functions.

These functions (except **tzset**) convert a time such as returned by **time(2)**.

<b><i>FUNCTION</i></b>	<b><i>REFERENCE</i></b>	<b><i>BRIEF DESCRIPTION</i></b>
<b>asctime</b>	<b>ctime(3C)</b>	Return string representation of date and time.
<b>ctime</b>	<b>ctime(3C)</b>	Return string representation of date and time, given integer form.
<b>gmtime</b>	<b>ctime(3C)</b>	Return Greenwich Mean Time.
<b>localtime</b>	<b>ctime(3C)</b>	Return local time.
<b>tzset</b>	<b>ctime(3C)</b>	Set time zone field from environment variable.

### **Miscellaneous Functions**

These functions support a wide variety of operations. Some of these are numerical conversion, password file and group file access, memory allocation, random number generation, and table management. These functions are automatically located and included in a program being compiled. No command line request is needed since these functions are part of the C library.

Some of these functions require declarations to be included. These are described following the descriptions of the functions.



**Numerical Conversion**

The following functions perform numerical conversion.

<i><b>FUNCTION</b></i>	<i><b>REFERENCE</b></i>	<i><b>BRIEF DESCRIPTION</b></i>
<b>a64l</b>	<b>a64l(3C)</b>	Convert string to base 64 ASCII.
<b>atof</b>	<b>atof(3C)</b>	Convert string to floating.
<b>atoi</b>	<b>atof(3C)</b>	Convert string to integer.
<b>atol</b>	<b>atof(3C)</b>	Convert string to long.
<b>frexp</b>	<b>frexp(3C)</b>	Split floating into mantissa and exponent.
<b>l3tol</b>	<b>l3tol(3C)</b>	Convert 3-byte integer to long.
<b>ltol3</b>	<b>l3tol(3C)</b>	Convert long to 3-byte integer.
<b>ldexp</b>	<b>frexp(3C)</b>	Combine mantissa and exponent.
<b>l64a</b>	<b>a64l(3C)</b>	Convert base 64 ASCII to string.
<b>modf</b>	<b>frexp(3C)</b>	Split mantissa into integer and fraction.

**DES Algorithm Access**

The following functions allow access to the Data Encryption Standard (DES) algorithm used on the UNIX operating system. The DES algorithm is implemented with variations to frustrate use of hardware implementations of the DES for key search.

<i>FUNCTION</i>	<i>REFERENCE</i>	<i>BRIEF DESCRIPTION</i>
<b>crypt</b>	<b>crypt(3C)</b>	Encode string.
<b>encrypt</b>	<b>crypt(3C)</b>	Encode/decode string of 0s and 1s.
<b>setkey</b>	<b>crypt(3C)</b>	Initialize for subsequent use of <b>encrypt</b> .

### Group File Access

The following functions are used to obtain entries from the group file. Declarations for these functions must be included in the program being compiled with the line:

```
#include <grp.h>
```

<i>FUNCTION</i>	<i>REFERENCE</i>	<i>BRIEF DESCRIPTION</i>
<b>endgrent</b>	<b>getgrent(3C)</b>	Close group file being processed.
<b>getgrent</b>	<b>getgrent(3C)</b>	Get next group file entry.
<b>getgrgid</b>	<b>getgrent(3C)</b>	Return next group with matching gid.
<b>getgrnam</b>	<b>getgrent(3C)</b>	Return next group with matching name.
<b>setgrent</b>	<b>getgrent(3C)</b>	Rewind group file being processed.
<b>fggetgrent</b>	<b>getgrent(3C)</b>	Get next group file entry from a specified file.

### Password File Access

These functions are used to search and access information stored in the password file (/etc/passwd). Some functions require declarations that can be included in the program being compiled by adding the line:

```
#include <pwd.h>
```

<i><b>FUNCTION</b></i>	<i><b>REFERENCE</b></i>	<i><b>BRIEF DESCRIPTION</b></i>
<b>endpwent</b>	<b>getpwent(3C)</b>	Close password file being processed.
<b>getpw</b>	<b>getpw(3C)</b>	Search password file for uid.
<b>getpwent</b>	<b>getpwent(3C)</b>	Get next password file entry.
<b>getpwnam</b>	<b>getpwent(3C)</b>	Return next entry with matching name.
<b>getpwuid</b>	<b>getpwent(3C)</b>	Return next entry with matching uid.
<b>putpwent</b>	<b>putpwent(3C)</b>	Write entry on stream.
<b>setpwent</b>	<b>getpwent(3C)</b>	Rewind password file being accessed.
<b>fgetpwent</b>	<b>getpwent(3C)</b>	Get next password file entry from a specified file.

### Parameter Access

The following functions provide access to several different types of parameters. None require any declarations.

<b><i>FUNCTION</i></b>	<b><i>REFERENCE</i></b>	<b><i>BRIEF DESCRIPTION</i></b>
<b>getopt</b>	<b>getopt(3C)</b>	Get next option from option list.
<b>getcwd</b>	<b>getcwd(3C)</b>	Return string representation of current working directory.
<b>getenv</b>	<b>getenv(3C)</b>	Return string value associated with environment variable.
<b>getpass</b>	<b>getpass(3C)</b>	Read string from terminal without echoing.
<b>putenv</b>	<b>putenv(3C)</b>	Change or add value of an environment variable.

### **Hash Table Management**

The following functions are used to manage hash search tables. The header file associated with these functions should be included in the program being compiled. This is accomplished by including the line:

```
#include <search.h>
```

near the beginning of any file using the search functions.

<b><i>FUNCTION</i></b>	<b><i>REFERENCE</i></b>	<b><i>BRIEF DESCRIPTION</i></b>
<b>hcreate</b>	<b>hsearch(3C)</b>	Create hash table.
<b>hdestroy</b>	<b>hsearch(3C)</b>	Destroy hash table.
<b>hsearch</b>	<b>hsearch(3C)</b>	Search hash table for entry.

## Binary Tree Management

The following functions are used to manage a binary tree. The header file associated with these functions should be included in the program being compiled. This is accomplished by including the line:

```
#include <search.h>
```

near the beginning of any file using the search functions.

<i>FUNCTION</i>	<i>REFERENCE</i>	<i>BRIEF DESCRIPTION</i>
<b>tdelete</b>	<b>tsearch(3C)</b>	Deletes nodes from binary tree.
<b>tfind</b>	<b>tsearch(3C)</b>	Find element in binary tree.
<b>tsearch</b>	<b>tsearch(3C)</b>	Look for and add element to binary tree.
<b>twalk</b>	<b>tsearch(3C)</b>	Walk binary tree.

## Table Management

The following functions are used to manage a table. Since none of these functions allocate storage, sufficient memory must be allocated before using these functions. The header file associated with these functions should be included in the program being compiled. This is accomplished by including the line:

```
#include <search.h>
```

near the beginning of any file using the search functions.

## C LIBRARIES

<i>FUNCTION</i>	<i>REFERENCE</i>	<i>BRIEF DESCRIPTION</i>
<b>bsearch</b>	<b>bsearch(3C)</b>	Search table using binary search.
<b>lfind</b>	<b>lsearch(3C)</b>	Find element in library tree.
<b>lsearch</b>	<b>lsearch(3C)</b>	Look for and add element in binary tree.
<b>qsort</b>	<b>qsort(3C)</b>	Sort table using quick-sort algorithm.

### Memory Allocation

The following functions provide a means by which memory can be dynamically allocated or freed.

<i>FUNCTION</i>	<i>REFERENCE</i>	<i>BRIEF DESCRIPTION</i>
<b>calloc</b>	<b>malloc(3C)</b>	Allocate zeroed storage.
<b>free</b>	<b>malloc(3C)</b>	Free previously allocated storage.
<b>malloc</b>	<b>malloc(3C)</b>	Allocate storage.
<b>realloc</b>	<b>malloc(3C)</b>	Change size of allocated storage.

The following is another set of memory allocation functions available.

<i>FUNCTION</i>	<i>REFERENCE</i>	<i>BRIEF DESCRIPTION</i>
<b>calloc</b>	<b>malloc(3X)</b>	Allocate zeroed storage.
<b>free</b>	<b>malloc(3X)</b>	Free previously allocated storage.

<b>malloc</b>	<b>malloc(3X)</b>	Allocate storage.
<b>mallopt</b>	<b>malloc(3X)</b>	Control allocation algorithm.
<b>mallinfo</b>	<b>malloc(3X)</b>	Space usage.
<b>realloc</b>	<b>malloc(3X)</b>	Change size of allocated storage.

### **Pseudorandom Number Generation**

The following functions are used to generate pseudorandom numbers. The functions that end with **48** are a family of interfaces to a pseudorandom number generator based upon the linear congruent algorithm and 48-bit integer arithmetic. The **rand** and **srand** functions provide an interface to a multiplicative congruential random number generator with period of 232.

<b><i>FUNCTION</i></b>	<b><i>REFERENCE</i></b>	<b><i>BRIEF DESCRIPTION</i></b>
<b>drand48</b>	<b>drand48(3C)</b>	Random double over the interval [0 to 1).
<b>lcong48</b>	<b>drand48(3C)</b>	Set parameters for <b>drand48</b> , <b>lrand48</b> , and <b>mrnd48</b> .
<b>lrand48</b>	<b>drand48(3C)</b>	Random long over the interval [0 to $2^{31}$ ).
<b>mrnd48</b>	<b>drand48(3C)</b>	Random long over the interval [ $-2^{31}$ to $2^{31}$ ).
<b>rand</b>	<b>rand(3C)</b>	Random integer over the interval [0 to 32767).

0

0.

0



## Chapter 18

# THE OBJECT AND MATH LIBRARIES

C-18

### GENERAL

This chapter describes the Object and Math Libraries that are supported on the UNIX operating system. A library is a collection of related functions and/or declarations that simplify programming effort. All of the functions described are also described in section 3, 3M, 3X and 3F of the Software Development System manual. Most of the declarations described are in section 5. The three main libraries on the UNIX system are:

- |                     |                                                                                                                                                                                                                                                                                                   |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>C library</b>    | This is the basic library for C language programs. The C library is composed of functions and declarations used for file access, string testing and manipulation, character testing and manipulation, memory allocation, and other functions. This library is described in the preceding chapter. |
| <b>Object file</b>  | This library provides functions for the access and manipulation of object files. This library is described later in this chapter.                                                                                                                                                                 |
| <b>Math library</b> | This library provides exponential, bessel functions, logarithmic, hyperbolic, and trigonometric functions. This library is also described later in this chapter.                                                                                                                                  |

## THE OBJECT FILE LIBRARY

The object file library provides functions for the access and manipulation of object files. Some functions locate portions of an object file such as the symbol table, the file header, sections, and line number entries associated with a function. Other functions read these types of entries into memory. For a description of the format of an object file, see "The Common Object File Format" in the Software Development System manual.

This library consists of several portions. The functions reside in */usr/lib/libld.a* and are located and loaded during the compiling of a C language program by a command line request. The form of this request is:

```
cc file -lld
```

which causes the link editor to search the object file library. The argument **-lld** must appear after all files that reference functions in *libld.aR*.

*In addition, various header files must be included. This is accomplished by including the line:*

```
#include <stdio.h>
#include <a.out.h>
#include <ldfcn.h>
```

<b>FUNCTION</b>	<b>REFERENCE</b>	<b>BRIEF DESCRIPTION</b>
<b>ldaclose</b>	<b>ldclose(3X)</b>	Close object file being processed.
<b>ldahread</b>	<b>ldahread(3X)</b>	Read archive header.
<b>ldaopen</b>	<b>ldopen(3X)</b>	Open object file for reading.
<b>ldclose</b>	<b>ldclose(3X)</b>	Close object file being processed.

<b>ldfhead</b>	<b>ldfhead(3X)</b>	Read file header of object file being processed.
<b>ldgetname</b>	<b>ldgetname(3X)</b>	Retrieve the name of an object file symbol table entry.
<b>ldlinit</b>	<b>ldlread(3X)</b>	Prepare object file for reading line number entries via <b>ldlitem</b> .
<b>ldlitem</b>	<b>ldlread(3X)</b>	Read line number entry from object file after <b>ldlinit</b> .
<b>ldlread</b>	<b>ldlread(3X)</b>	Read line number entry from object file.
<b>ldlseek</b>	<b>ldlseek(3X)</b>	Seeks to the line number entries of the object file being processed.
<b>ldnlseek</b>	<b>ldlseek(3X)</b>	Seeks to the line number entries of the object file being processed given the name of a section.
<b>ldnrseek</b>	<b>ldrseek(3X)</b>	Seeks to the relocation entries of the object file being processed given the name of a section.
<b>ldnshread</b>	<b>ldshread(3X)</b>	Read section header of the named section of the object file being processed.
<b>ldnsseek</b>	<b>ldsseek(3X)</b>	Seeks to the section of the object file being processed given the name of a section.

<b>ldohseek</b>	<b>ldohseek(3X)</b>	Seeks to the optional file header of the object file being processed.
<b>ldopen</b>	<b>ldopen(3X)</b>	Open object file for reading.
<b>ldrseek</b>	<b>ldrseek(3X)</b>	Seeks to the relocation entries of the object file being processed.
<b>ldshread</b>	<b>ldshread(3X)</b>	Read section header of an object file being processed.
<b>ldsseek</b>	<b>ldsseek(3X)</b>	Seeks to the section of the object file being processed.
<b>ldtbindex</b>	<b>ldtbindex(3X)</b>	Returns the long index of the symbol table entry at the current position of the object file being processed.
<b>ldtbread</b>	<b>ldtbread(3X)</b>	Reads a specific symbol table entry of the object file being processed.
<b>ldtbseek</b>	<b>ldtbseek(3X)</b>	Seeks to the symbol table of the object file being processed.
<b>sgetl</b>	<b>sputl(3X)</b>	Access long integer data in a machine independant format.
<b>sputl</b>	<b>sputl(3X)</b>	Translate a long integer into a machine independant format.

### Common Object File Interface Macros (ldfcn.h)

The interface between the calling program and the object file access routines is based on the defined type **LDFILE** which is defined in the header file **ldfcn.h** (see **ldfcn(4)**). The primary purpose of this structure is to provide uniform access to both simple object files and to object files that are members of an archive file.

The function **ldopen(3X)** allocates and initializes the **LDFILE** structure and returns a pointer to the structure to the calling program. The fields of the **LDFILE** structure may be accessed individually through the following macros: the **type** macro returns the magic number of the file, which is used to distinguish between archive files and simple object files. The **IOPTR** macro returns the file pointer which was opened by **ldopen(3X)** and is used by the input/output functions of the C library. The **OFFSET** macro returns the file address of the beginning of the object file. This value is non-zero only if the object file is a member of the archive file. The **HEADER** macro accesses the file header structure of the object file.

Additional macros are provided to access an object file. These macros parallel the input/output functions in the C library; each macro translates a reference to an **LDFILE** structure into a reference to its file descriptor field. The available macros are described in **ldfcn(4)**.

## THE MATH LIBRARY

The math library consists of functions and a header file. The functions are located and loaded during the compiling of a C language program by a command line request. The form of this request is:

```
cc file -lm
```

which causes the link editor to search the math library. In addition to the request to load the functions, the header file of the math

## THE OBJECT AND MATH LIBRARIES

library should be included in the program being compiled. This is accomplished by including the line:

```
#include <math.h>
```

near the beginning of the (first) file being compiled.

The functions are grouped into the following categories:

- Trigonometric functions
- Bessel functions
- Hyperbolic functions
- Miscellaneous functions.

### Trigonometric Functions

These functions are used to compute angles (in radian measure), sines, cosines, and tangents. All of these values are expressed in double precision.

<i>FUNCTION</i>	<i>REFERENCE</i>	<i>BRIEF DESCRIPTION</i>
<b>acos</b>	<b>trig(3M)</b>	Return arc cosine.
<b>asin</b>	<b>trig(3M)</b>	Return arc sine.
<b>atan</b>	<b>trig(3M)</b>	Return arc tangent.
<b>atan2</b>	<b>trig(3M)</b>	Return arc tangent of a ratio.
<b>cos</b>	<b>trig(3M)</b>	Return cosine.

<b>sin</b>	<b>trig(3M)</b>	Return sine.
<b>tan</b>	<b>trig(3M)</b>	Return tangent.

### Bessel Functions

These functions calculate bessel functions of the first and second kinds of several orders for real values. The bessel functions are **j0**, **j1**, **jn**, **y0**, **y1**, and **yn**. The functions are located in section **bessel(3M)**.

### Hyperbolic Functions

These functions are used to compute the hyperbolic sine, cosine, and tangent for real values.

<b>FUNCTION</b>	<b>REFERENCE</b>	<b>BRIEF DESCRIPTION</b>
<b>cosh</b>	<b>sinh(3M)</b>	Return hyperbolic cosine.
<b>sinh</b>	<b>sinh(3M)</b>	Return hyperbolic sine.
<b>tanh</b>	<b>sinh(3M)</b>	Return hyperbolic tangent.

### Miscellaneous Functions

These functions cover a wide variety of operations, such as natural logarithm, exponential, and absolute value. In addition, several are provided to truncate the integer portion of double precision numbers.

<b>FUNCTION</b>	<b>REFERENCE</b>	<b>BRIEF DESCRIPTION</b>
<b>ceil</b>	<b>floor(3M)</b>	Returns the smallest integer not less than a given value.
<b>exp</b>	<b>exp(3M)</b>	Returns the exponential function of a given value.

## THE OBJECT AND MATH LIBRARIES

<b>fabs</b>	<b>floor(3M)</b>	Returns the absolute value of a given value.
<b>floor</b>	<b>floor(3M)</b>	Returns the largest integer not greater than a given value.
<b>fmod</b>	<b>floor(3M)</b>	Returns the remainder produced by the division of two given values.
<b>gamma</b>	<b>gamma(3M)</b>	Returns the natural log of the absolute value of the result of applying the gamma function to a given value.
<b>hypot</b>	<b>hypot(3M)</b>	Return the square root of the sum of the squares of two numbers.
<b>log</b>	<b>exp(3M)</b>	Returns the natural logarithm of a given value.
<b>log10</b>	<b>exp(3M)</b>	Returns the logarithm base ten of a given value.
<b>matherr</b>	<b>matherr(3M)</b>	Error-handling function.
<b>pow</b>	<b>exp(3M)</b>	Returns the result of a given value raised to another given value.
<b>sqrt</b>	<b>exp(3M)</b>	Returns the square root of a given value.



## Chapter 19

# LEXICAL ANALYZER GENERATOR (LEX)

C-19

### GENERAL

The **Lex** is a program generator that produces a program in a general purpose language that recognizes regular expressions. It is designed for lexical processing of character input streams. It accepts a high-level, problem oriented specification for character string matching. The regular expressions are specified by you (the user) in the source specifications given to **Lex**. The **Lex** program generator source is a table of regular expressions and corresponding program fragments. The table is translated to a program that reads an input stream, copies the input stream to an output stream, and partitions the input into strings that match the given expressions. As each such string is recognized, the corresponding program fragment is executed. The recognition of the expressions is performed by a deterministic finite automaton generated by **Lex**. The program fragments written by you are executed in the order in which the corresponding regular expressions occur in the input stream.

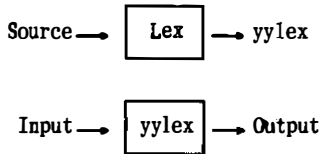
The user supplies the additional code beyond expression matching needed to complete the tasks, possibly including codes written by other generators. The program that recognizes the expressions is generated in the general purpose programming language employed for your program fragments. Thus, a high-level expression language is provided to write the string expressions to be matched while your freedom to write actions is unimpaired.

The **Lex** written code 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 user. Compatible run-time libraries for the different host languages are also provided. This makes **Lex** adaptable to different environments and different

# LEX

users. Each application may be directed to the combination of hardware and host language appropriate to the task, the user's background, and the properties of local implementations. At present, the only supported host language is the C language, although Fortran (in the form of Ratfor) has been available in the past. The **Lex** generator exists on the UNIX operating system, but the codes generated by **Lex** may be taken anywhere the appropriate compilers exist.

The **Lex** program generator turns the user's expressions and actions (called **source**) into the host general purpose language; the generated program is named **yylex**. The **yylex** program recognizes expressions in a stream (called **input**) and performs the specified actions for each expression as it is detected. See Figure 11-1.



**Figure 11-1. Overview of Lex**

For 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. This rule contains a regular expression that matches one or more instances of the characters

blank or tab (written for visibility, in accordance with the C language convention) and occurs 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," as in **QED**. No action is specified, so the program generated by **Lex yylex()** ignores these characters. Everything else is copied. To change any remaining string of blanks or tabs to a single blank, add another rule.

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

The coded instructions (generated for this source) scan for both rules at once, observe (at the termination of the string of blanks or tabs) whether or not there is a newline character, and then execute the desired rule action. The first rule matches all strings of blanks or tabs at the end of lines, and the second rule matches all remaining strings of blanks or tabs.

The **Lex** program generator can be used alone for simple transformations or for analysis and statistics gathering on a lexical level. The **Lex** generator can also be used with a parser generator to perform the lexical analysis phase; it is particularly easy to interface **Lex** and **yacc**. The **Lex** program recognizes only regular expressions; **yacc** writes parsers that accept a large class of context free grammars but requires a lower level analyzer to recognize input tokens. Thus, a combination of **Lex** and **yacc** is often appropriate. When used as a preprocessor for a later parser generator, **Lex** is used to partition the input stream; and the parser generator assigns structure to the resulting pieces. The flow of control in such a case is shown in Figure 11.2. Additional programs, written by other generators or by hand, can be added easily to programs written by **Lex**. You 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.

In the program written by **Lex**, the user'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 user to insert either

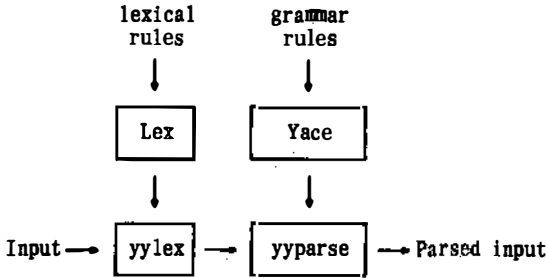


Figure 11-2. Lex With Yacc

declarations or additional statements in the routine containing the actions or to add subroutines outside this action routine.

The **Lex** program generator is not limited to a source that can be interpreted on the basis of one character look-ahead. For example, if there are two rules, one looking for "ab" and another for "abcdefg" and the input stream is "abcdefh," **Lex** recognizes "ab" and leaves the input pointer just before "cd ...". Such backup is more costly than the processing of simpler languages.

## LEX SOURCE

The general format of **Lex** source is

```

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

```

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

```
%%
```

(no definitions, no rules) which translates into a program that copies the input to the output unchanged.

In the outline of **Lex** programs shown above, the rules represent your control decisions. They are in a table containing

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

Thus an individual rule might be

```
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 language library function **printf** is used to print the string. The end of the expression is indicated by the first blank or tab character. If the action is merely a single C language 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 more useful example, suppose you desire to change a number of words from British to American spelling.

## LEX

The **Lex** rules such as:

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

would be a start. These rules are not sufficient since the word "petroleum" would become "gaseum".

## LEX REGULAR EXPRESSIONS

The definitions of regular expressions are very similar to those in **QED**. 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; 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 should 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 equivalent to the one above. Thus, by quoting every nonalphanumeric character being used as a text character, the user 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 a backslash (`\`) 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 language 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 except blank, tab, newline, and the list of operator characters 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 lowercase 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 uppercase letters, both lowercase letters, or both digits is implementation dependent and gets a warning message (e.g., [0-z] in ASCII is many more characters than 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 to indicate that the resulting string is complemented with respect to the computer character set. Thus:

```
[^abc]
```

matches all characters except "a", "b", or "c", including all special or control characters; or

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

is any character that is not a letter. The \ character provides the usual escapes within character class brackets.

**Arbitrary Character**

To match almost any character, the operator character (dot)

is the class of all characters except newline. Escaping into octal is possible although nonportable.

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



matches all printable ASCII characters 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 `+`. For example,

`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 lowercase 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

The **Lex** program recognizes a small amount of surrounding context. The two simplest operators for this are `^` and `$`. If the first character of an expression is `^`, the expression is only matched at the beginning of a line (after a newline character or at the beginning of the input stream). This never conflicts with the other meaning of `^` (complementation of character classes) since that only applies within the `[]` operators. If the very last character is `$`, the expression is only 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 followed by “cd”. Thus:

`ab$`

is the same as

ab/\n

Left context is handled in **Lex** by “start conditions” as explained later. 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 later.

### 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.

## LEX ACTIONS

When an expression written as above is matched, **Lex** executes the corresponding action. This part describes some features of **Lex** that aid in writing actions. Note that there is a default action that consists of copying the input to the output. This is performed on all strings not otherwise matched. Thus, the **Lex** user 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 that merely copies can be omitted. Also, a character combination that is omitted from the rules and that 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 language null statement, `;` as an action causes this result. A frequent rule is

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

which causes the three spacing characters (blank, tab, and newline) to be ignored.

Another easy way to avoid writing actions is the action character `!` which indicates that the action for this rule is the action for the next rule. The previous example could also have been written

```
" " !
" \\t" !
" \\n" ;
```

with the same result although in different style. The quotes around `\\n` and `\\t` are not required.

In more complex actions, you may often want to know the actual text that matched some expression like `"[a-z]+"`. The **Lex** program leaves this text in an external character array. Thus, to print the name found, a rule like

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

prints the string in *yytext*[]. The C language 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*[]. This 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 that is not desired. For example, if there is a rule that 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.

Sometimes it is more convenient to know the end of what has been found; hence, **Lex** also provides a count *yy leng* of the number of characters matched. To count both the number of words and the number of characters in words in the input, write

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

which accumulates in *chars* the number of characters in the words recognized. The last character in the string matched can be accessed by

```
yytext[yy leng-1]
```

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, *yy more*() 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, *yy less*(*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 in *yytext* to be retained. Further characters previously matched are returned to the input. This provides the same sort of look ahead offered by the / operator but in a different form.

Example:

Consider a language that 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[yyleng-1] == '\\')
        yymore();
    else
        ... normal user processing
}
```

will, when faced with a string such as "abc\def", first match the five characters "abc\"; then the call to *yymore()* will cause the next part of the string "def" to be tacked on the end. Note that the final quote terminating the string should be picked up in the code labeled "normal processing".

The function *yyless()* might be used to reprocess text in various circumstances. Consider the C language problem of distinguishing the ambiguity of "= - a ". Suppose it is desired to treat this as "= - a" but also to print a message: a rule might be

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

```

C-19

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 is 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 as follows:

1. *input()* returns the next input character.
2. *output(c)* writes the character “c” on the output.
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 user 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 cause input or output to be transmitted to or from strange places including other programs or internal memory. The character set used must be consistent in all routines and a value of zero returned by *input* must mean end of file. The relationship between *unput* and *input* must be retained or the **Lex** look ahead will not work. The **Lex** program does not look ahead at all if it does not have to, but every rule ending in +, \*, ?, or \$ or containing / implies look ahead. Look ahead is also necessary to match an expression that is a prefix of another expression. The standard **Lex** library imposes a 100-character limit on backup.

Another **Lex** library routine that you may 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 wrap-up on end of input. Sometimes, however, it is convenient to arrange for more input to arrive from a new source. In this case, the user 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 that recognizes end of file; the only access to this condition is through *yywrap*. In fact, unless a private version of *input()* is supplied, a file containing nulls cannot be handled since a value of 0 returned by *input* is taken to be end of file.

## **AMBIGUOUS SOURCE RULES**

The **Lex** program 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 that matched the same number of characters, the rule given first is preferred.



Thus, suppose the rules

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

are to be given in that order. If the input is "integers", it is taken as an identifier because

"[a-z]+"

matches eight characters while "integer" matches only seven. If the input is "integer", both rules match seven characters; and the keyword rule is selected because it was given first. Anything shorter (e.g., "int") does 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 appear to be a good way of recognizing a string in single quotes. However, 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 will match

```
'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 dot (.) operator

does not match `newline`. Thus expressions like `*stop-on-the-current-line`. Do not try to defeat this with expressions like `[\n]+` or equivalents; the **Lex** generated program tries 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 dot (.) does not include newline. Since “she” includes “he”, **Lex** normally *does not* recognize the instances of “he” included in “she” since once it has passed a “she” those characters are gone.

Sometimes the user desires to override this choice. The action *REJECT* means “go do the next alternative”. It causes whatever rule was second choice after the current rule to be executed. The position of the input pointer is adjusted accordingly. Suppose you really want to count the included instances of “he”. Use the following rule to change the previous example to accomplish the task.

```
she  {s++; REJECT;}
he   {h++; REJECT;}
\n   |
.    ;
```

After counting each expression, it is rejected; whenever appropriate, the other expression is then counted. In this example, you could note that “she” includes “he” but not vice versa and omit the *REJECT* action on “he”. In other cases, it is not possible to state which input characters are 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 then 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 2-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.

The action *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* executed, the user must not have used *unput* to change the characters forthcoming from the input stream. This is the only restriction on the user’s ability to manipulate the not-yet-processed input.

## LEX SOURCE DEFINITIONS

Recalling the format of the **Lex** source,

```
{definitions}
%%
{rules}
%%
{user routines}
```

So far, only the rules have been described. You need additional options to define variables for use in the program and for use by **Lex**. Variables can go either in the definitions section or in the rules section.

Remember **Lex** is generating 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 not part of a **Lex** rule or action that 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.

Lines that begin with a blank or tab and that contain a comment are passed through to the generated program. This can be used to include comments in either the **Lex** source or the generated code; the comments should follow the host language convention.

2. Anything included between lines containing only %{ 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 causes the string given as a translation to be associated with the name. The name and translation must be separated by at least one blank or tab, and the name must begin with a letter. The translation can then be called out by the `{name}` syntax in a rule. Using `{D}` for the digits and `{E}` for an exponent field, for example, 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. 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

## LEX

itself for larger source programs. These possibilities are discussed later.

### USAGE

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 language standard library.

On the UNIX operating system, the library is accessed by the loader flag **-ll**. So an appropriate set of commands is

```
lex source
cc lex.yy.c -ll
```

The resulting program is placed on the usual file *a.out* for later execution. To use **Lex** with **yacc**, see part "LEX AND YACC". Although the default **Lex** I/O routines use the C language standard library, the **Lex** automata themselves do not do so; if private versions of *input*, *output*, and *unput* are given, the library is avoided.

### LEX AND YACC

To use **Lex** with **yacc**, observe that **Lex** writes a program named *yylex()* (the name required by **yacc** for its analyzer). Normally, the default main program on 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 ends with

```
return(token);
```

where the appropriate token value is returned. 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. If the grammar is to be named "good" and the lexical rules are to be named "better", the UNIX software command sequence could be

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

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

## EXAMPLES

As a problem, consider copying an input file while adding three to every positive number divisible by seven. A suitable **Lex** source program follows:

```
%%
    int k;
[0-9]+ {
    k = atoi(yytext);
    if (k%7 == 0)
        printf(" %d", k+3);
    else
        printf(" %d", k);
}
```

The rule "[0-9]+" recognizes strings of digits; *atoi()* converts the digits to binary and stores the result in "k". The operator % (remainder) is used to check whether "k" is divisible by seven; if it is, "k" is incremented by three as it is written out. It may be objected that this program alters such input items as "49.63" or "X7". Furthermore, it increments the absolute value of all negative numbers divisible by seven. To avoid this, 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 dot (.) or preceded by a letter will be picked up by one of the last two rules and not changed. The “if-else” has been replaced by a C language 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 that histograms the lengths of words, where a word is defined as a string of letters:

```

                                int lengs[100];
%%
[a-z]+          lengs[yyval]++;
.              |
\n            ;
%%
yywrap()
{
int i;
printf(" Length No. words\n");
for(i=0; i<100; i++)
    if (lengs[i] > 0)
        printf(" %5d%10d\n",i,lengs[i]);
return(1);
}

```

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 wrap-up. If *yywrap* returns zero (false), it implies that further input is available and the program is to continue reading and processing. Providing a *yywrap* (that never returns true) causes an infinite loop.



## 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 part 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 that recognize the need to change the environment in which the following input text is analyzed and that set a parameter to reflect the change. This may be a flag explicitly tested by the user's action code; this 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 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 began with the letter "a", changing "magic" to "second" on every line which began with the letter "b", and changing "magic" to "third" on every line which began 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 <> brackets;

```
<name1>expression
```

is a rule that 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;
```

resets 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 as follows:

```
%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 user's code.

## 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 used to return values in *ytext()*. 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 that translate the characters, **Lex** must be given a translation table that is in the definitions section and must be bracketed by lines containing only

**%T**; the translation table contains lines of the form

```
{integer} {character string}
```

which indicate the value associated with each character.

## **SUMMARY OF SOURCE FORMAT**

The general form of a **Lex** source file is

```
{definitions}  
% %  
{rules}  
% %  
{user 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 conditions 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 “a” selects the parameter as follows:

<b>Letter</b>	<b>Parameter</b>
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 operators shown in Figure 11-3.

OPERATOR	DESCRIPTION
x	the character "x".
" x"	an "x", even if x is an operator.
\x	an "x", even if x is an operator.
[xy]	the character x or y.
[x-z]	the characters x, y, or z.
[^x]	any character but x.
.	any character but newline.
^x	an x at the beginning of a line.
<y>x	an x when Lex is in start condition y.
x\$	an x at the end of a line.
x?	an optional x.
x*	0,1,2, ... instances of x.
x+	1,2,3, ... instances of x.
x y	an x or a y.
(x)	an x.
x/y	an x but only if followed by y.
{xx}	the translation of xx from the definitions section.
x{m,n}	m through n occurrences of x.

Figure 11-3. Operators and Descriptions

## CAVEATS AND BUGS

There are pathological expressions that produce exponential growth of the tables when converted to deterministic machines; 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* executed, the user must not have used *unput* to change the characters forthcoming from the input stream. This is the only restriction on the user's ability to manipulate the not-yet-processed input.

## Chapter 20

# YET ANOTHER COMPILER-COMPILER (yacc)

### GENERAL

The **yacc** program provides a general tool for imposing structure on the input to a computer program. The **yacc** user 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. The **yacc** program then generates a function to control the input process. This function, called a *parser*, calls the user-supplied low-level input routine (the *lexical analyzer*) to pick up the basic items (called *tokens*) from the input stream. These tokens are organized according to the input structure rules, called *grammar rules*. When one of these rules has been recognized, then user 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.

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

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

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

where “date”, “month\_name”, “day”, and “year” represent structures of interest in the input process; presumably, “month name”, “day”, and “year” are defined elsewhere. The comma is enclosed in single quotes. This implies that the comma is to appear literally in the input. The colon and semicolon merely serve as punctuation in the

rule and have no significance in controlling the input. With proper definitions, the input

July 4, 1776

might be matched by the rule.

An important part of the input process is carried out by the lexical analyzer. This user routine reads the input stream, recognizes the lower-level structures, and communicates these tokens to the parser. For historical reasons, a structure recognized by the lexical analyzer is called a "terminal symbol", while the structure recognized by the parser is called a "nonterminal symbol". To avoid confusion, terminal symbols will usually be 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 only needs to recognize individual letters, and "month name" is a nonterminal symbol. Such low-level rules tend to waste time and space and may complicate the specification beyond the ability of **yacc** to deal with it. Usually, the lexical analyzer recognizes the month names and returns an indication that a "month name" is seen. In this case, "month name" is a "token".

Literal characters such as a comma 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

on input. In most cases, this new rule could be "slipped in" to a working system with minimal effort and little danger of disrupting existing input.

The input being read 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 **yacc** program has been extensively used in numerous practical applications, including **lint**, the Portable C Compiler, and a system for typesetting mathematics.

The remainder of this document describes the following subjects as they relate to **yacc**

- Basic process of preparing a **yacc** specification
- Parser operation
- Handling ambiguities
- Handling operator precedences in arithmetic expressions
- Error detection and recovery
- The operating environment and special features of the parsers **yacc** produces
- Suggestions to improve the style and efficiency of the specifications
- Advanced topics.

In addition, there are four appendices. Appendix 12.1 is a brief example, and Appendix 12.2 is a summary of the **yacc** input syntax. Appendix 12.3 gives an example using some of the more advanced features of **yacc**, and Appendix 12.4 describes mechanisms and syntax no longer actively supported but provided for historical continuity with older versions of **yacc**.

## BASIC SPECIFICATIONS

Names refer to either tokens or nonterminal symbols. The **yacc** program requires token names to be declared as such. In addition, it is often desirable to include the lexical analyzer as part of the specification file. It may be useful to include other programs as well. Thus, every specification file consists of three sections: the *declarations*, (*grammar*) *rules*, and *programs*. The sections are separated by double percent (%%) marks. (The percent symbol is generally used in **yacc** specifications as an escape character.)

In other words, a full specification file looks like

```
declarations
%%
rules
%%
programs
```

when each section is used.

The declaration section may be empty, and if the programs section is omitted, the second %% mark may also be omitted. The smallest legal yacc specification is

```
%%
rules
```

since the other two sections may be omitted.

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

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

```
A : BODY ;
```

where "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, dots, underscores, and noninitial digits. Uppercase and lowercase letters are distinct. The names used in the body of a grammar rule may represent tokens or nonterminal symbols.

## YACC

A literal consists of a character enclosed in single quotes ('). As in C language, the backslash (\) is an escape character within literals, and all the C language escapes are recognized. Thus:

```
'\n'  newline
'\r'  return
'\''  single quote ( ' )
'\\'  backslash ( \ )
'\t'  tab
'\b'  backspace
'\f'  form feed
'\xxx' "xxx" in octal
```

are understood by yacc. 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
  ;
```

by using the vertical bar. 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 by

```
empty : ;
```

which is understood by `yacc`.

Names representing tokens must be declared. This is most simply done by writing

```
%token name1 name2 ...
```

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

Of all the nonterminal symbols, 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 desirable to declare the start symbol explicitly in the declarations section using the `%start` keyword

```
%start symbol
```

to define the start symbol.

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

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

## ACTIONS

With each grammar rule, the user 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 language statement and as such can do input and output, call subprograms, and alter external vectors and variables. An action is specified by one or more statements enclosed in curly braces ({} and {} ). For example:

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

and

```
XXX : YYY 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, the action

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

does nothing but return the value of one.

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. If the rule is

$$A : B C D ;$$

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

The rule

$$\text{expr} : '(' \text{ expr } ')';$$

provides a more concrete example. The value returned by this rule is usually the value of the "expr" in parentheses. This can be indicated by

$$\begin{aligned} \text{expr} : & '(' \text{ expr } ') \\ & \{ \\ & \quad \$$ = \$2; \\ & \} \end{aligned}$$

By default, the value of a rule is the value of the first element in it ( $\$1$ ). 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 rules. Sometimes, it is desirable to get control before a rule is fully parsed. The **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. The **yacc** program actually treats the above example as if it had been written

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

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

where  $\$ACT$  is an empty action.

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



maintain the tree structure desired. 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. Then parse tree can be built by supplying actions such as

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

in the specification.

The user 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 with `yy`. The user should avoid such names.

In these examples, all the values are integers. A discussion of values of other types is found in the part "ADVANCED TOPICS".

## LEXICAL ANALYSIS

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

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 the user. In either case, the **#define** mechanism of C language 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 );
        ...
    }
    ...
}

```

to return the appropriate token.

The intent is to return a 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 is defined as the token number associated with the token DIGIT.

This mechanism leads to clear, easily modified lexical analyzers. The only pitfall to avoid is using any token names in the grammar that are reserved or significant in C language or the parser. For example, the use of token names **if** or **while** 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.

As mentioned above, the token numbers may be chosen by **yacc** or the user. 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 end-marker must have token number 0 or negative. This token number cannot be redefined by the user. Thus, all lexical analyzers should be prepared to return 0 or a negative number as a token upon reaching the end of their input.

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

## PARSER OPERATION

The **yacc** program turns the specification file into a C language program, which parses the input according to the specification given. The algorithm used to go from the specification to the parser is complex and will not be discussed here. The parser itself, however, is relatively simple and understanding how it works will 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 is also capable of reading and remembering the next input token (called the *look-ahead* token). The *current state* is always the one on the top of the stack. The states of the finite state machine are given small integer labels. Initially, the machine is in state 0 (the stack contains only state 0) and no look-ahead token has been read.

The machine has only four actions available—*shift*, *reduce*, *accept*, and *error*. A step of the parser is done as follows:

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

The *shift* action is the most common action the parser takes. Whenever a shift action is taken, there is always a look-ahead token. For example, in state 56 there may be an action

```
IF shift 34
```

which says, in state 56, if the look-ahead token is IF, the current state (56) is pushed down on the stack, and state 34 becomes the current state (on the top of the stack). The look-ahead token is cleared.

The *reduce* action keeps the stack from growing without bounds. Reduce actions are appropriate when the parser has seen the right-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 look-ahead token to decide whether to reduce or not (usually it is not necessary). In fact, the default action (represented by a dot) is often a reduce action.

Reduce actions are associated with individual grammar rules. Grammar rules are also given small integer numbers, and this leads to some confusion. The action

. reduce 18

refers to grammar rule 18, while the action

IF shift 34

refers to state 34.

Suppose the rule

A : x y z ;

is being reduced. 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 look-ahead token is cleared by a shift but is not affected by a *goto*. In any case, the uncovered state contains an entry such as

A goto 20

causing state 20 to be pushed onto the stack and become 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 of the stacks. The uncovered state is in fact the current state.

The reduce action is also important in the treatment of user-supplied actions and values. When a rule is reduced, the code supplied with the rule is executed before the stack is adjusted. In addition to the stack holding the states, another stack running in parallel with it holds the values returned from the lexical analyzer and the actions. When a shift takes place, the external variable "yyval" is copied onto the value stack. After the return from the user code, the reduction is carried out. When the *goto* action is done, the external variable "yyval" is copied onto the value stack. The pseudo-variables \$1, \$2, etc., refer to the value stack.

The other two parser actions are conceptually much simpler. The *accept* action indicates that the entire input has been seen and that it matches the specification. This action appears only when the look-ahead token is the end-marker and indicates that the parser 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 look-ahead token) cannot be followed by anything that would result in a legal input. The parser reports an error and attempts to recover the situation and resume parsing. The error recovery (as opposed to the detection of error) will be discussed later.

Consider:

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

as a yacc specification.

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
```

where the actions for each state are specified and 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. The following input

```
DING DONG DELL
```

can be used to track the operations of the parser. 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 and becomes the look-ahead token. The action in state 0 on *DING* is *shift 3*, state 3 is pushed onto the stack, and the look-ahead token is cleared. State 3 becomes the current state. The next token, *DONG*, is read and becomes the look-ahead token. The action in state 3 on the token *DONG* is *shift 6*, state 6 is pushed onto the stack, and the look-ahead is cleared. The stack now contains 0, 3, and 6. In state 6, without even consulting the look-ahead, the parser reduces by

```
sound : DING DONG
```

which is rule 2. Two states, 6 and 3, are popped off of the stack uncovering state 0. Consulting the description of state 0 (looking for a goto on *sound*),

```
sound goto 2
```

is obtained. State 2 is pushed onto the stack and becomes the current state.



In state 2, the next token, *DELL*, must be read. The action is *shift 5*, so state 5 is pushed onto the stack, which now has 0, 2, and 5 on it, and the look-ahead token is cleared. In state 5, the only action is to reduce by rule 3. This has one symbol on the right-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 and the end-marker is obtained indicated by *\$end* in the *y.output* file. The action in state 1 (when the end-marker is seen) successfully ends the parse.

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

## AMBIGUITY AND CONFLICTS

A set of grammar rules is ambiguous if there is some input string that can be structured in two or more different ways. For example, the grammar rule

$$\text{expr} : \text{expr} \text{'-'} \text{expr}$$

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 completely specify the way that all complex inputs should be structured. For example, if the input is

$$\text{expr} - \text{expr} - \text{expr}$$

the rule allows this input to be structured as either

$$( \text{expr} - \text{expr} ) - \text{expr}$$

or as

$$\text{expr} - ( \text{expr} - \text{expr} )$$

(The first is called “left association”, the second “right association”).

The **yacc** program detects such ambiguities when it is attempting to build the parser. Given the input

$$\text{expr} - \text{expr} - \text{expr}$$

consider the problem that confronts the parser. When the parser has read the second **expr**, the input seen

$$\text{expr} - \text{expr}$$

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

$$- \text{expr}$$

and again reduce. The effect of this is to take the left associative interpretation.

Alternatively, if the parser sees

$$\text{expr} - \text{expr}$$

it could defer the immediate application of the rule and continue reading the input until

$$\text{expr} - \text{expr} - \text{expr}$$

is seen. It could then apply the rule to the rightmost three symbols reducing them to "expr" which results in

expr - expr

being left. 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 one of two legal things, a shift or a reduction. It has no way of deciding between them. This is called a "shift/reduce conflict". It 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 the choice to make in a given situation is called a "disambiguating rule".

The **yacc** program 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 in favor of shifts when there is a choice. Rule 2 gives the user rather crude control over the behavior of the parser in this situation, but reduce/reduce conflicts should be avoided when possible.

Conflicts 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

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

which is a fragment from a programming language involving an “if-then-else” statement. In these rules, “IF” and “ELSE” are tokens, “cond” is a nonterminal symbol describing conditional (logical) expressions, and “stat” is a nonterminal symbol describing statements. The first rule 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 ( C1 ) IF ( C2 ) S1 ELSE S2
```

can be structured according to these rules in two ways

```

IF ( C1 )
{
    IF ( C2 )
        S1
}
ELSE
    S2

```

or

```

IF ( C1 )
{
    IF ( C2 )
        S1
    ELSE
        S2
}

```

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

and is looking at the "ELSE". It can immediately reduce by the simple-if rule to get

```
IF ( C1 ) stat
```

and then read the remaining input

```
ELSE S2
```

and reduce

```
IF ( C1 ) stat ELSE S2
```

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, "S2" read, and then the right-hand portion of

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

can be reduced by the if-else rule to get

```
IF ( C1 ) 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, such as

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

have already been seen. 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
```

where the first line describes the conflict—giving the state and the input symbol. The ordinary state description gives the grammar rules active in the state and the parser actions. Recall that the underline marks the portion of the grammar rules 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. In state 23, the alternative action [describing a dot (.)] is to be done if the input symbol is not mentioned explicitly in the actions. In this case, if the input symbol is not "ELSE", the parser reduces to

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

by grammar rule 18.

Once again, notice that the numbers following "shift" commands refer to other states, while the numbers following "reduce" commands refer to grammar rule numbers. In the *y.output* file, the rule numbers are printed after those rules which can be reduced. In most one states, there is reduce action possible in the state and this is the default command. The user who encounters unexpected shift/reduce conflicts will probably want to look at the verbose output to decide whether the default actions are appropriate.

## PRECEDENCE

---

There is one common situation where the rules given above for resolving conflicts are not sufficient. This is in the parsing of arithmetic expressions. Most of the commonly used constructions for arithmetic expressions can be naturally described by the notion of precedence levels for operators, together with information about left or right associativity. It turns out that ambiguous grammars with appropriate disambiguating rules can be used to create parsers that are faster and easier to write than parsers constructed from unambiguous grammars. The basic notion is to write grammar rules of the 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 user specifies the precedence or binding strength of all the operators and the associativity of the binary operators. This information is sufficient to allow **yacc** to resolve the parsing conflicts in accordance with these rules and construct a parser that realizes the desired precedences and associativities.

The precedences and associativities are attached to tokens in the declarations section. This is done by a series of lines beginning with 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 operator `.LT.` 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) ) )
```

in order to perform the correct precedence of operators. 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. The keyword **%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 causes the precedence of the grammar rule to become that of the following token name or literal. For example, the rules

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

%%

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

might be used to give unary minus the same precedence as multiplication.

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 "cookbook" fashion until some experience has been gained. The `y.out.put` file is very useful in deciding whether the parser is actually doing what was intended.

## ERROR HANDLING

Error handling is an extremely difficult area, and many of the problems are semantic ones. When an error is found, for example, it 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 user some control over this process, `yacc` provides a simple, but reasonably general feature. The token name “error” is reserved for error handling. This name can be used in grammar rules. In effect, it suggests places where errors are expected and recovery might take place. The parser pops its stack until it enters a state where the token “error” is legal. It then behaves as if the token “error” were the current look-ahead token and performs the action encountered. The look-ahead token is then reset to the token that caused the error. If no special error rules 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
```

means that on a syntax error the parser attempts to skip over the statement in which the error is seen. More precisely, the parser scans 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. Rules such as

```
stat : error ';
```

are somewhat easier. Here, when there is an error, the parser attempts to skip over the statement but does so by skipping to the next semicolon. All tokens after the error and before the next semicolon cannot be shifted and are discarded. When the semicolon is seen, this rule will be reduced and any "cleanup" action associated with it performed.

C-20

Another form of error rule arises in interactive applications where it may be desirable to permit a line to be reentered after an error. The following example

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

is one way to do this. 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 force the parser to believe that error recovery has been accomplished. The statement

```
yerror ;
```

in an action resets the parser to its normal mode. The last example can be rewritten as

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

which is somewhat better.

As previously mentioned, the token seen immediately after the "error" symbol is the input token at which the error was discovered. Sometimes, this is inappropriate; for example, an error recovery action might take upon itself the job of finding the correct place to resume input. In this case, the previous look-ahead token must be cleared. The 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 user) that attempted to advance the input to the beginning of the next valid statement. After this routine is called, the next token returned by *yylex* is presumably the first token in a legal statement. The old illegal token must be discarded and the error state reset. A rule similar to

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

could perform this.

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

## THE “yacc” ENVIRONMENT

When the user inputs a specification to **yacc**, the output is a file of C language programs, called *y.tab.c* on most systems. (Due to local file system conventions, the names may differ from installation to installation.) The function produced by **yacc** is called *yyparse()*; it is an integer valued function. When it is called, it in turn repeatedly calls *yylex()*, the lexical analyzer supplied by the user (see “LEXICAL ANALYSIS”), to obtain input tokens. Eventually, an error is detected, *yyparse()* returns the value 1, and no error recovery is possible, or the lexical analyzer returns the end-marker token and the parser accepts. In this case, *yyparse()* returns the value 0.

The user must provide a certain amount of environment for this parser in order to obtain a working program. For example, as with every C language 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.

These two routines must be supplied in one form or another by the user. To ease the initial effort of using **yacc**, a library has been provided with default versions of *main()* and *yyerror()*. The name of this library is system dependent; on many systems, the library is accessed by a `-ly` argument to the loader. The source codes

```
main()
{
    return ( yyparse() );
}
```

and

---

```
# include <stdio.h>

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

show the triviality of these default programs. The argument to *yyerror()* is a string containing an error message, usually the string "syntax error". The average application wants to do better than this. Ordinarily, the program should keep track of the input line number and print it along with the message when a syntax error is detected. The external integer variable *yychar* contains the look-ahead token number at the time the error was detected. This may be of some interest in giving better diagnostics. Since the *main()* program is probably supplied by the user (to read arguments, etc.), the *yacc* library is useful only in small projects or in the earliest stages of larger ones.

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

## HINTS FOR PREPARING SPECIFICATIONS

This part 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 are a few style hints.

1. Use all uppercase letters for token names and all lowercase letters for nonterminal names. This rule comes under the heading of "knowing who to blame when things go wrong".
2. Put grammar rules and actions on separate lines. This allows either to be changed without an automatic need to change the other.
3. 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.
4. 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 easily added.
5. Indent rule bodies by two tab stops and action bodies by three tab stops.

The example in Appendix 12.1 is written following this style, as are the examples in this section (where space permits). The user 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 ;
```

match this algorithm. These rules such as

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

and

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

frequently arise when writing specifications of sequences and lists. 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 is a bit bigger; and the items are seen and reduced from right to left. More seriously, an internal stack in the parser is in danger of overflowing if a very long sequence is read. Thus, the user should use left recursion wherever reasonable.

It is worth considering if a sequence with zero elements has any meaning, and if so, consider writing the sequence specification as

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

using an empty rule. 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,

```

% {
    int dflag;
% }
... other declarations ...

%%

prog : decls stats
    ;

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

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

... other rules ...

```

specifies a program that consists of zero or more declarations followed by zero or more statements. 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 “back-door” 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 you 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 user 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*, i.e., forbidden for use as variable names. There are powerful stylistic reasons for preferring this.

## ADVANCED TOPICS

This part 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**. The **YYACCEPT** macro causes *yyparse()* to return the value 0; **YYERROR** causes the parser to behave as if the current input symbol had been a syntax error; *yyperror()* is called, and error recovery takes place. These mechanisms can be used to simulate parsers with multiple end-markers or context sensitive syntax checking.

## Accessing Values in Enclosing Rules

An action 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.

```

sent : adj noun verb adj noun
    {
        look at the sentence ...
    }
;
adj  : THE
    {
        $$ = THE;
    }
    | YOUNG
    {
        $$ = YOUNG;
    }
;
noun : DOG
    {
        $$ = DOG;
    }
    | CRONE
    {
        if( $0 == YOUNG )
        {
            printf( " what?\n" );
        }
        $$ = CRONE;
    }
;
...

```

In this case, the digit may be 0 or negative. In the action following the word CRONE, a check is made that the preceding token shifted was not YOUNG. Obviously, this is only possible when a great deal is known about what might precede the symbol "noun" in the input. There is also a distinctly unstructured flavor about this. Nevertheless, at times this mechanism prevents a great deal of trouble especially when a few combinations are to be excluded from an otherwise regular structure.

## Support for Arbitrary Value Types

~~By default, the values returned by actions and the lexical analyzer~~ are integers. The `yacc` program 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 is strictly type checked. The `yacc` value stack is declared to be a *union* of the various types of values desired. The user declares the union and associates union member names to each token and nonterminal symbol having a value. When the value is referenced through a `$$` or `$n` construction, `yacc` will automatically insert the appropriate union name so that no unwanted conversions take place. In addition, type checking commands such as `lint` is 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 user 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 nonterminals. 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 user includes

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

in the declaration section. This declares the `yacc` value stack and the external variables `yyval` 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 have said

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

instead. 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> '+' '-'
```

causes any reference to values returned by these two tokens to be tagged with the union member name *optype*. Another keyword, `%type`, is used to associate union member names with nonterminals. Thus, one might say

```
%type <nodetype> expr stat
```

to associate the union member *nodetype* with the nonterminal symbols "expr" and "stat".

There remains a couple of cases where these mechanisms are insufficient. If there is an action within a rule, the value returned by this action has no *a priori* type. Similarly, reference to left context values (such as `$0`) leaves `yacc` with no easy way of knowing the type. In this case, a type can be imposed on the reference by

inserting a union member name between < and > immediately after the first \$. The example

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

shows this usage. This syntax has little to recommend it, but the situation arises rarely.

A sample specification is given in Appendix 12.3. 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.

## APPENDIX 12.1

### A Simple Example

This example gives the complete yacc applications for a small desk calculator; the calculator has 26 registers labeled "a" through "z" and accepts arithmetic expressions made up of the operators +, -, \*, /, % (mod operator), & (bitwise and), | (bit wise or), and assignments. If an expression at the top level is an assignment, the value is printed; otherwise, the expression is printed. As in C language, an integer that begins with 0 (zero) is assumed to be octal; otherwise, it is assumed to be decimal.

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



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

```

% {
# includes<stdio.h>
# includes<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 rule section */

list    : /* empty */
        | list stat '\n'
        | list error '\n'

        {
            yyerrork;
        }
        ;

stat    : expr
        {
            printf( " %dn", $1 );
        }
        | LETTER '=' expr
        {
            regs[$1] = $3
        }
        }

```

```

;
expr  : '(' expr ')'
      {
        $$ = $2;
      }
      | expr '+' expr
      {
        $$ = $1 + $3
      }
      | expr '-' expr
      {
        $$ = $1 - $3
      }
      | expr '*' expr
      {
        $$ = $1 * $3;
      }
      | expr '/' expr
      {
        $$ = $1/$3;
      }
      | exp '%' expr
      {
        $$ = $1 % $3
      }
      | expr '&' expr
      {
        $$ = $1 & $3;
      }
      | expr '|' expr
      {
        $$ = $1 | $3
      }
      | '-' expr %prec UMINUS
      {
        $$ = - $2;
      }
      | LETTER
      {
        $$ = reg[$1];
      }

```

```

        | number
        ;

number  : DIGIT
        {
            $$ = $1; base = ($1==0) ? 8 ; 10;
        }
        | number DIGIT
        {
            $$ = bas * $1 + $2
        }
        ;

%% /* start of program */

yylex() /* lexical analysis routine */
{
    /* return LETTER for lowercase letter,
       yyval = 0 through 25*/

    /* returns DIGIT for digit, yyval = 0 through 9*/
    /* all other characters are returned immediately */

    int c;
                                /*skip blanks*/
    while (c=getchar( ) == " )
        ;

    /* c is now nonblank */

    if( islower( c ))
    {
        yyval = c- 'a';
        return( LETTER );
    }
    if( isdigit( c ))
    {
        yyval = c-'0';
        return( DIGIT );
    }
    return( c );
}

```

## APPENDIX 12.2

## YACC Input Syntax

This appendix has a description of 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 decides whether the next token (skipping blanks, newlines, and 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\_IDENTIFIERS.

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

```
/* basic entries */
```

```
%token IDENTIFIER /* includes identifiers and literals */
```

```
%token C_IDENTIFIER /* identifier (but not literal)
    followed by a colon */
```

```
%token NUMBER /* [0-9]+ */
```

```
/* reserved words: %type=> TYPE %left=>LEFT,etc. */
```

```
%token LEFT RIGHT NONASSOC TOKEN PREC TYPE START UNION
```

```
%token MARK /* the %% mark */
```

```
%token LCURL /* the % { mark */
```

```
%token RCURL /* the % } mark */
```

```
/* ASCII character literals stand for themselves */
```

```
%token spec
```

```
%%
```

```

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
      ;

defs : START IDENTIFIER
      | UNION
      {
          Copy union definition to output
      }
      | LCURL
      {
          Copy C code to output file
          RCURL
      }
      | ndefs rword tag nlist
      ;

rword : TOKEN
      | LEFT
      | RIGHT
      | NONASSOC
      | TYPE
      ;

tag : /* empty: union tag is optional */
     | '<' IDENTIFIER '>'
     ;

nlist : nmno
       | nlist nmno
       | nlist ',' nmno

```

```

;

nmno : IDENTIFIER /*Note: literal illegal with % type */
      | IDENTIFIER NUMBER /* Note: illegal with % type */
;

/* rule section */

rules : C_IDENTIFIER rbody proc
       | rules rule
;

rule : C_IDENTIFIER rbody prec
      | ' ' rbody prec
;

rbody : /* empty */
       | rbody IDENTIFIER
       | rbody act
;

act : '{'
     {
       Copy action translate $$' etc.
     }
     '}'
;

Bprec : /* empty */
       | PREC IDENTIFIER
       | PREC IDENTIFIER act
       | prec';
;

```

## APPENDIX 12.3

### An Advanced Example

This appendix gives an example of a grammar using some of the advanced features. The desk calculator example in Appendix 12.1 is modified to provide a desk calculator that does floating point interval arithmetic. The calculator understands floating point constants; the arithmetic operations  $+$ ,  $-$ ,  $*$ ,  $/$ , unary  $-$  "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 Appendix 12.1; assignments returns no value and prints nothing while expressions print the (floating or interval) value.

This example explores a number of interesting features of **yacc** and C language. Intervals are represented by a structure consisting of the left and right endpoint values stored as doubles. This structure is given a type name, **INTERVAL**, by using *typedef*. The **yacc** value stack can also contain floating point scalars and integers (used to index into the arrays holding the variable values). Notice that the entire strategy depends strongly on being able to assign structures and unions in C language. 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. 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 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 value expression in the second example, but this fact is not known until the comma is read. By this time,, 2.5 is finished, and the parser cannot go back and change its mind. More generally, it might be necessary to look ahead an arbitrary number of tokens to decide whether to convert a scalar to an interval. This problem is evaded by having two rules for each binary interval valued operator—one when the left operand is a scalar and one when the left operand is an interval. In the second case, the right operand must be an interval, so the conversion 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 conflict 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 language library routine *atof()* is used to do the actual conversion from a character string to a double precision value. If the lexical analyzer detects an error, it responds by returning a token that is illegal in the grammar provoking a syntax error in the parser and thence error recovery.

```
% {
```

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



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

```
INTERVAL vmul(), vdiv( );
```

```
double atof( );
```

```
double dreg[ 26 ];
INTERVAL vreg[ 26 ];
```

```
% }
```

```
%start line
```

```
%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 )0,$1.1o,$1.1hi );
    }
| DREG '=' '\n'
    {
        dreg[$1] = $3;
    }
| VREG '=' vexp '\n'
    {
        vreg[$1] = $3;
    }
| error '\n'
    {
        yyerrork;
    }
;

dexp : CONST
    | DREG
    {
        $$ = dreg[$1]
    }
| dexp '+' dexp
    {
        $$ = $1 + $3
    }

```

```

| dexp '-' dexp
{
    $$ = $1 - $3
}
| dexp '*' dexp
{
    $$ = $1 * $3
}
| dexp '/' dexp
{
    $$ = $1 / $3
}
}
| '-' dexp %prec UMINUS
{
    $$ = - $2
}
| '(' dexp ')'
{
    $$ = $2
}
}
;
vexpp : dexp
{
    $$hi = $$lo = $1;
}
| '(' dexp', dexp' )
{
    $$lo = $2;
    $$hi = $4;
    If( $$lo > $$hi )
    {

```

```

        printf( " interval out of order n" );
        YYERROR;
    }
}
| VREG
{
    $$ = vreg[$1]
}
| vexp '+' vexp
{
    $$hi = $1hi + $3hi;
    $$lo = $1lo + $3lo
}
| dexp '+' vexp
{
    $$hi = $1 + $3hi;
    $$lo = $1 + $3lo
}
| vexp '=' vexp
{
    $$hi = $1hi - $3lo;
    $$lo = $1lo - $3hi
}
| dexp '-' vdep
{
    $$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;
        $$ = vdiv( $1.lo, $1.hi, $3 )
    }
    | '-' vexp %prec UMINUS
    {
        $$ .hi = -$2.lo; $$ .lo = -$2.hi
    }
    | '(' vexp ')'
    {
        $$ = $2
    }
    ;
%%

#define BSZ 50 /* buffer size for floating point number */

/* lexical analysis */

yylex( )
{
    register c;

                                /* skip over blanks */
    ;

```

```

if( isupper( c ) )
{
    yylval.ival = c - 'A'
    return( VREG );
}
if( islower( c ) )
{

    yylval.ival = c - 'a',
    return( DREG );
}

/* gobble up digits. points, exponents */
if( idigit( c ) || c=='.' )
{

    char buf[BSZ+1], *cp = buf;
    int dot = 0, exp = 0;

    for( ; (cp-buf)<BSZ ; ++cp,c=getchar( ) )
    {

        *cp = c;
        if( isdigit( c ) )
            continue;
        if( c == '.' )
        {
            if( dot++ || exp )
                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-buff) >= 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 */ routine */

```

```

    INTERVAL v;

```

```

    if( a>b )

```

```

    {

```

```

        v.hi = a;

```

```

        v.lo = b;

```

```

    }

```

```

    else

```

```

    {

```

```

        v.hi = b;

```

```

        v.lo = a;

```

```

    }

```

```

    if( c>d )

```

```

    {

```

```

        if( c>v.hi )

```

```

            v.hi = c;

```

```

        if( d<v.lo )

```

```

            v.lo = d;

```

```

    }

```

```

    else

```

```

    {

```

```

        if( d>v.hi )

```

```

            v.hi = d;

```

```

        if( c<v.lo )

```

```

            v.lo = c;

```

```

    }

```

```

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

```

## APPENDIX 12.4

### Old Features Supported But Not Encouraged

This appendix mentions synonyms and features that 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 literal.

The use of multicharacter 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. Action may also have the form

`= { ... }`

and the curly braces can be dropped if the action is a single C language statement.

6. The C language code between `%{` and `%}` use to be permitted at the head of the rules section as well as in the declaration section.



## TABLE OF CONTENTS OF COMMANDS

<b>2. System Calls and Error Numbers</b>	
intro.....	introduction to system calls and error numbers
access.....	determine accessibility of a file
acct.....	enable or disable process accounting
alarm.....	set a process alarm clock
brk.....	change data segment space allocation
chdir.....	change working directory
chmod.....	change mode of file
chown.....	change owner and group of a file
chroot.....	change root directory
close.....	close a file descriptor
creat.....	create a new file or rewrite an existing one
dup.....	duplicate an open file descriptor
exec.....	execute a file
exit.....	terminate process
fcntl.....	file control
fork.....	create a new process
getpid.....	get process, process group, and parent process IDs
getuid.....	get real user, effective user, real group, and effective group IDs
ioctl.....	control device
kill.....	send a signal to a process or a group of processes
link.....	link to a file
lseek.....	move read/write file pointer
mknod.....	make a directory, or a serial or ordinary file
mount.....	mount a file system
msgctl.....	message control operations
msgget.....	get message queue
msgop.....	message operations
nice.....	change priority of a process
open.....	open for reading or writing
pause.....	suspend process until signal
pipe.....	create an interprocess channel
lock.....	lock process, text, or data in memory
profil.....	execution time profile
prace.....	process trace
read.....	read from file
semctl.....	semaphore control operations
semget.....	get set of semaphores
semop.....	semaphore operations
setpgrp.....	set process group ID
setuid.....	set user and group IDs
shmctl.....	shared memory control operations
shmunget.....	get shared memory segment
shmop.....	shared memory operations
signal.....	specify what to do on receipt of a signal
stat.....	get file status
stime.....	set time
sync.....	update super-block
time.....	get time
times.....	get process and child process times
uadmin.....	administrative control
ulimit.....	get and set user limits
umask.....	set and get file creation mask
umount.....	unmount a file system
uname.....	get name of current UNIX system
unlink.....	remove directory entry
ustat.....	get file system statistics
utime.....	set file access and modification times
wait.....	wait for child process to stop or terminate
write.....	write on a file



## NAME

intro - introduction to system calls and error numbers

## SYNOPSIS

```
#include <errno.h>
```

## DESCRIPTION

This section describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible returned value. This is almost always -1; the individual descriptions specify the details. An error number is also made available in the external variable *errno*. *Errno* is not cleared on successful calls, so it should be tested only after an error has been indicated.

Each system call description attempts to list all possible error numbers. The following is a complete list of the error numbers and their names as defined in *<errno.h>*.

- 1 EPERM Not owner  
Typically this error indicates an attempt to modify a file in some way forbidden except to its owner or super-user. It is also returned for attempts by ordinary users to do things allowed only to the super-user.
- 2 ENOENT No such file or directory  
This error occurs when a file name is specified and the file should exist but doesn't, or when one of the directories in a path name does not exist.
- 3 ESRCH No such process  
No process can be found corresponding to that specified by *pid* in *kill* or *ptrace*.
- 4 EINTR Interrupted system call  
An asynchronous signal (such as interrupt or quit), which the user has elected to catch, occurred during a system call. If execution is resumed after processing the signal, it will appear as if the interrupted system call returned this error condition.
- 5 EIO I/O error  
Some physical I/O error has occurred. This error may in some cases occur on a call following the one to which it actually applies.
- 6 ENXIO No such device or address  
I/O on a special file refers to a subdevice which does not exist, or beyond the limits of the device. It may also occur when, for example, a tape drive is not on-line or no disk pack is loaded on a drive.
- 7 E2BIG Arg list too long  
An argument list longer than 5,120 bytes is presented to a member of the *exec* family.
- 8 ENOEXEC Exec format error  
A request is made to execute a file which, although it has the appropriate permissions, does not start with a valid magic number [see *a.out(4)*].
- 9 EBADF Bad file number  
Either a file descriptor refers to no open file, or a read (respectively, write) request is made to a file which is open only for writing (respectively, reading).

## INTRO(2)

- 10 ECHILD No child processes  
A *wait* was executed by a process that had no existing or unwaited-for child processes.
- 11 EAGAIN No more processes  
A *fork* failed because the system's process table is full or the user is not allowed to create any more processes.
- 12 ENOMEM Not enough space  
During an *exec*, *brk*, or *sbrk*, a program asks for more space than the system is able to supply. This is not a temporary condition; the maximum space size is a system parameter. The error may also occur if the arrangement of text, data, and stack segments requires too many segmentation registers, or if there is not enough swap space during a *fork*.
- 13 EACCES Permission denied  
An attempt was made to access a file in a way forbidden by the protection system.
- 14 EFAULT Bad address  
The system encountered a hardware fault in attempting to use an argument of a system call.
- 15 ENOTBLK Block device required  
A non-block file was mentioned where a block device was required, e.g., in *mount*.
- 16 EBUSY Device or resource busy  
An attempt was made to mount a device that was already mounted or an attempt was made to dismount a device on which there is an active file (open file, current directory, mounted-on file, active text segment). It will also occur if an attempt is made to enable accounting when it is already enabled. The device or resource is currently unavailable.
- 17 EEXIST File exists  
An existing file was mentioned in an inappropriate context, e.g., *link*.
- 18 EXDEV Cross-device link  
A link to a file on another device was attempted.
- 19 ENODEV No such device  
An attempt was made to apply an inappropriate system call to a device; e.g., read a write-only device.
- 20 ENOTDIR Not a directory  
A non-directory was specified where a directory is required, for example in a path prefix or as an argument to *chdir(2)*.
- 21 EISDIR Is a directory  
An attempt was made to write on a directory.
- 22 EINVAL Invalid argument  
Some invalid argument (e.g., dismounting a non-mounted device; mentioning an undefined signal in *signal*, or *kill*; reading or writing a file for which *lseek* has generated a negative pointer). Also set by the math functions described in the (3M) entries of this manual.
- 23 ENFILE File table overflow  
The system file table is full, and temporarily no more *opens* can be accepted.

- 24 **EMFILE** Too many open files  
No process may have more than 20 file descriptors open at a time. When a record lock is being created with *fcntl*, there are too many files with record locks on them.
- 25 **ENOTTY** Not a character device  
An attempt was made to *ioctl(2)* a file that is not a special character device.
- 26 **ETXTBSY** Text file busy  
An attempt was made to execute a pure-procedure program that is currently open for writing. Also an attempt to open for writing a pure-procedure program that is being executed.
- 27 **EFBIG** File too large  
The size of a file exceeded the maximum file size (1,082,201,088 bytes) or *ULIMIT*; see *ulimit(2)*.
- 28 **ENOSPC** No space left on device  
During a *write* to an ordinary file, there is no free space left on the device. In *fcntl*, the setting or removing of record locks on a file cannot be accomplished because there are no more record entries left on the system.
- 29 **ESPIPE** Illegal seek  
An *lseek* was issued to a pipe.
- 30 **EROFS** Read-only file system  
An attempt to modify a file or directory was made on a device mounted read-only.
- 31 **EMLINK** Too many links  
An attempt to make more than the maximum number of links (1000) to a file.
- 32 **EPIPE** Broken pipe  
A write on a pipe for which there is no process to read the data. This condition normally generates a signal; the error is returned if the signal is ignored.
- 33 **EDOM** Math argument  
The argument of a function in the math package (3M) is out of the domain of the function.
- 34 **ERANGE** Result too large  
The value of a function in the math package (3M) is not representable within machine precision.
- 35 **ENOMSG** No message of desired type  
An attempt was made to receive a message of a type that does not exist on the specified message queue; see *msgop(2)*.
- 36 **EIDRM** Identifier Removed  
This error is returned to processes that resume execution due to the removal of an identifier from the file system's name space [see *msgctl(2)*, *semctl(2)*, and *shmctl(2)*].
- 45 **EDEADLK** Deadlock  
A deadlock situation was detected and avoided.

**DEFINITIONS****Process ID**

Each active process in the system is uniquely identified by a positive integer called a process ID. The range of this ID is from 1 to 30,000.

## INTRO(2)

### Parent Process ID

A new process is created by a currently active process; see *fork(2)*. The parent process ID of a process is the process ID of its creator.

### Process Group ID

Each active process is a member of a process group that is identified by a positive integer called the process group ID. This ID is the process ID of the group leader. This grouping permits the signaling of related processes; see *kill(2)*.

### Tty Group ID

Each active process can be a member of a terminal group that is identified by a positive integer called the tty group ID. This grouping is used to terminate a group of related processes upon termination of one of the processes in the group; see *exit(2)* and *signal(2)*.

### Real User ID and Real Group ID

Each user allowed on the system is identified by a positive integer called a real user ID.

Each user is also a member of a group. The group is identified by a positive integer called the real group ID.

An active process has a real user ID and real group ID that are set to the real user ID and real group ID, respectively, of the user responsible for the creation of the process.

### Effective User ID and Effective Group ID

An active process has an effective user ID and an effective group ID that are used to determine file access permissions (see below). The effective user ID and effective group ID are equal to the process's real user ID and real group ID respectively, unless the process or one of its ancestors evolved from a file that had the set-user-ID bit or set-group ID bit set; see *exec(2)*.

### Super-user

A process is recognized as a *super-user* process and is granted special privileges if its effective user ID is 0.

### Special Processes

The processes with a process ID of 0 and a process ID of 1 are special processes and are referred to as *proc0* and *proc1*.

*Proc0* is the scheduler. *Proc1* is the initialization process (*init*). *Proc1* is the ancestor of every other process in the system and is used to control the process structure.

### File Descriptor

A file descriptor is a small integer used to do I/O on a file. The value of a file descriptor is from 0 to 19. A process may have no more than 20 file descriptors (0-19) open simultaneously. A file descriptor is returned by system calls such as *open(2)*, or *pipe(2)*. The file descriptor is used as an argument by calls such as *read(2)*, *write(2)*, *ioctl(2)*, and *close(2)*.

### File Name

Names consisting of 1 to 14 characters may be used to name an ordinary file, special file, or directory.

These characters may be selected from the set of all character values in the range of octal values 1 through 0177 excluding the ASCII code of / (slash).

Note that it is generally unwise to use \*, ?, [, or ] as part of file names because of the special meaning attached to these characters by the shell. See *sh(1)*. Although permitted, it is advisable to avoid the use of



unprintable characters in file names.

### Path Name and Path Prefix

A path name is a null-terminated character string starting with an optional slash (/), followed by zero or more directory names separated by slashes, optionally followed by a file name.

More precisely, a path name is a null-terminated character string constructed as follows:

```
<path-name>::=<file-name>|<path-prefix><file-name>/
<path-prefix>::=<rtprefix>|/<rtprefix>
<rtprefix>::=<dirname>|<rtprefix><dirname>/
```

where <file-name> is a string of 1 to 14 characters other than the ASCII slash and null, and <dirname> is a string of 1 to 14 characters (other than the ASCII slash and null) that names a directory.

If a path name begins with a slash, the path search begins at the *root* directory. Otherwise, the search begins from the current working directory.

A slash by itself names the root directory.

Unless specifically stated otherwise, the null path name is treated as if it named a nonexistent file.

### Directory

Directory entries are called links. By convention, a directory contains at least two links, . and .., referred to as *dot* and *dot-dot* respectively. Dot refers to the directory itself and dot-dot refers to its parent directory.

### Root Directory and Current Working Directory

Each process has associated with it a concept of a root directory and a current working directory for the purpose of resolving path name searches. The root directory of a process need not be the root directory of the root file system.

### File Access Permissions

Read, write, and execute/search permissions on a file are granted to a process if one or more of the following are true:

The effective user ID of the process is super-user.

The effective user ID of the process matches the user ID of the owner of the file and the appropriate access bit of the "owner" portion (0700) of the file mode is set.

The effective user ID of the process does not match the user ID of the owner of the file, and the effective group ID of the process matches the group of the file and the appropriate access bit of the "group" portion (070) of the file mode is set.

The effective user ID of the process does not match the user ID of the owner of the file; the effective group ID of the process does not match the group ID of the file; the appropriate access bit of the "other" portion (07) of the file mode is set.

Otherwise, the corresponding permissions are denied.

### Message Queue Identifier

A message queue identifier (msqid) is a unique positive integer created by a *msgget*(2) system call. Each msqid has a message queue and a data structure associated with it.

## INTRO(2)

The data structure is referred to as *msqid\_ds* and contains the following members:

```
struct ipc_perm msg_perm; /* operation permission struct */
ushort msg_qnum;          /* number of msgs on q */
ushort msg_qbytes;        /* max number of bytes on q */
ushort msg_lspid;         /* pid of last msgsnd operation */
ushort msg_lrpid;         /* pid of last msgrcv operation */
time_t msg_stime;         /* last msgsnd time */
time_t msg_rtime;         /* last msgrcv time */
time_t msg_ctime;         /* last change time */
                          /* Times measured in secs since */
                          /* 00:00:00 GMT, Jan. 1, 1970 */
```

**Msg\_perm** is an *ipc\_perm* structure that specifies the message operation permission (see below). This structure includes the following members:

```
ushort cuid;             /* creator user id */
ushort cgid;             /* creator group id */
ushort uid;              /* user id */
ushort gid;              /* group id */
ushort mode;            /* r/w permission */
```

**Msg\_qnum** is the number of messages currently on the queue. **Msg\_qbytes** is the maximum number of bytes allowed on the queue. **Msg\_lspid** is the process id of the last process that performed a *msgsnd* operation. **Msg\_lrpid** is the process id of the last process that performed a *msgrcv* operation. **Msg\_stime** is the time of the last *msgsnd* operation, **msg\_rtime** is the time of the last *msgrcv* operation, and **msg\_ctime** is the time of the last *msgctl(2)* operation that changed a member of the above structure.

### Message Operation Permissions

In the *msgop(2)* and *msgctl(2)* system call descriptions, the permission required for an operation is given as "{token}", where "token" is the type of permission needed interpreted as follows:

00400	Read by user
00200	Write by user
00060	Read, Write by group
00006	Read, Write by others

Read and Write permissions on a *msqid* are granted to a process if one or more of the following are true:

The effective user ID of the process is super-user.

The effective user ID of the process matches **msg\_perm.[c]uid** in the data structure associated with *msqid* and the appropriate bit of the "user" portion (0600) of **msg\_perm.mode** is set.

The effective user ID of the process does not match **msg\_perm.[c]uid** and the effective group ID of the process matches **msg\_perm.[c]gid** and the appropriate bit of the "group" portion (060) of **msg\_perm.mode** is set.

The effective user ID of the process does not match **msg\_perm.[c]uid** and the effective group ID of the process does not match **msg\_perm.[c]gid** and the appropriate bit of the "other" portion (06) of **msg\_perm.mode** is set.

Otherwise, the corresponding permissions are denied.

### Semaphore Identifier

A semaphore identifier (*semid*) is a unique positive integer created by a *semget(2)* system call. Each *semid* has a set of semaphores and a data structure associated with it. The data structure is referred to as *semid\_ds* and contains the following members:

```

struct ipc_perm sem_perm; /* operation permission struct */
ushort sem_nsems;        /* number of sems in set */
time_t sem_otime;       /* last operation time */
time_t sem_ctime;       /* last change time */
                        /* Times measured in secs since */
                        /* 00:00:00 GMT, Jan. 1, 1970 */

```

**sem\_perm** is an *ipc\_perm* structure that specifies the semaphore operation permission (see below). This structure includes the following members:

```

ushort cuid;            /* creator user id */
ushort cgid;            /* creator group id */
ushort uid;             /* user id */
ushort gid;             /* group id */
ushort mode;           /* r/a permission */

```

The value of **sem\_nsems** is equal to the number of semaphores in the set. Each semaphore in the set is referenced by a positive integer referred to as a *sem\_num*. *sem\_num* values run sequentially from 0 to the value of *sem\_nsems* minus 1. **sem\_otime** is the time of the last *semop(2)* operation, and **sem\_ctime** is the time of the last *semctl(2)* operation that changed a member of the above structure.

A semaphore is a data structure that contains the following members:

```

ushort semval;          /* semaphore value */
short sempid;           /* pid of last operation */
ushort semncnt;         /* # awaiting semval > cval */
ushort semzcnt;         /* # awaiting semval = 0 */

```

**semval** is a non-negative integer. **Sempid** is equal to the process ID of the last process that performed a semaphore operation on this semaphore. **Semncnt** is a count of the number of processes that are currently suspended awaiting this semaphore's *semval* to become greater than its current value. **Semzcnt** is a count of the number of processes that are currently suspended awaiting this semaphore's *semval* to become zero.

### Semaphore Operation Permissions

In the *semop(2)* and *semctl(2)* system call descriptions, the permission required for an operation is given as "{token}", where "token" is the type of permission needed interpreted as follows:

00400	Read by user
00200	Alter by user
00060	Read, Alter by group
00006	Read, Alter by others

Read and Alter permissions on a *semid* are granted to a process if one or more of the following are true:

The effective user ID of the process is super-user.

The effective user ID of the process matches **sem\_perm.cuid** in the data structure associated with *semid* and the appropriate bit of the "user" portion (0600) of **sem\_perm.mode** is set.

The effective user ID of the process does not match **sem\_perm.cuid** and the effective group ID of the process

matches `sem_perm.cjgid` and the appropriate bit of the "group" portion (060) of `sem_perm.mode` is set.

The effective user ID of the process does not match `sem_perm.cjuid` and the effective group ID of the process does not match `sem_perm.cjgid` and the appropriate bit of the "other" portion (06) of `sem_perm.mode` is set.

Otherwise, the corresponding permissions are denied.

### Shared Memory Identifier

A shared memory identifier (`shmid`) is a unique positive integer created by a `shmget(2)` system call. Each `shmid` has a segment of memory (referred to as a shared memory segment) and a data structure associated with it. The data structure is referred to as `shmid_ds` and contains the following members:

```

struct ipc_perm shm_perm; /* operation permission struct */
int shm_segsz; /* size of segment */
ushort shm_cpid; /* creator pid */
ushort shm_lpid; /* pid of last operation */
short shm_nattch; /* number of current attaches */
time_t shm_atime; /* last attach time */
time_t shm_dtime; /* last detach time */
time_t shm_ctime; /* last change time */
/* Times measured in secs since */
/* 00:00:00 GMT, Jan. 1, 1970 */

```

`Shm_perm` is an `ipc_perm` structure that specifies the shared memory operation permission (see below). This structure includes the following members:

```

ushort cuid; /* creator user id */
ushort cgid; /* creator group id */
ushort uid; /* user id */
ushort gid; /* group id */
ushort mode; /* r/w permission */

```

`Shm_segsz` specifies the size of the shared memory segment. `Shm_cpid` is the process id of the process that created the shared memory identifier. `Shm_lpid` is the process id of the last process that performed a `shmop(2)` operation. `Shm_nattch` is the number of processes that currently have this segment attached. `Shm_atime` is the time of the last `shmat` operation, `shm_dtime` is the time of the last `shmdt` operation, and `shm_ctime` is the time of the last `shmctl(2)` operation that changed one of the members of the above structure.

### Shared Memory Operation Permissions

In the `shmop(2)` and `shmctl(2)` system call descriptions, the permission required for an operation is given as "{token}", where "token" is the type of permission needed interpreted as follows:

00400	Read by user
00200	Write by user
00060	Read, Write by group
00006	Read, Write by others

Read and Write permissions on a `shmid` are granted to a process if one or more of the following are true:

The effective user ID of the process is super-user.

The effective user ID of the process matches `shm_perm.cjuid` in the data structure associated with `shmid` and the appropriate bit of the "user" portion (0600) of `shm_perm.mode` is set.

The effective user ID of the process does not match **shm\_perm.[c]uid** and the effective group ID of the process matches **shm\_perm.[c]gid** and the appropriate bit of the "group" portion (060) of **shm\_perm.mode** is set.

The effective user ID of the process does not match **shm\_perm.[c]uid** and the effective group ID of the process does not match **shm\_perm.[c]gid** and the appropriate bit of the "other" portion (06) of **shm\_perm.mode** is set.

Otherwise, the corresponding permissions are denied.

#### Shared Memory Operations in the Small Memory Model

The shared memory system calls *shmget(2)*, *shmop(2)*, and *shmctl(2)* are not allowed in System V/286 small model programs, since the small memory model allows only one data segment.

#### SEE ALSO

*close(2)*, *ioctl(2)*, *open(2)*, *pipe(2)*, *read(2)*, *write(2)*, *intro(3)*.

**This page intentionally left blank.**

**NAME**

access — determine accessibility of a file

**SYNOPSIS**

```
int access (path, amode)
char *path;
int amode;
```

**DESCRIPTION**

*Path* points to a path name naming a file. *Access* checks the named file for accessibility according to the bit pattern contained in *amode*, using the real user ID in place of the effective user ID and the real group ID in place of the effective group ID. The bit pattern contained in *amode* is constructed as follows:

04	read
02	write
01	execute (search)
00	check existence of file

Access to the file is denied if one or more of the following are true:

[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	Read, write, or execute (search) permission is requested for a null path name.
[ENOENT]	The named file does not exist.
[EACCES]	Search permission is denied on a component of the path prefix.
[EROFS]	Write access is requested for a file on a read-only file system.
[ETXTBSY]	Write access is requested for a pure procedure (shared text) file that is being executed.
[EACCESS]	Permission bits of the file mode do not permit the requested access.
[EFAULT]	<i>Path</i> points outside the allocated address space for the process.

The owner of a file has permission checked with respect to the "owner" read, write, and execute mode bits. Members of the file's group other than the owner have permissions checked with respect to the "group" mode bits, and all others have permissions checked with respect to the "other" mode bits.

**RETURN VALUE**

If the requested access is permitted, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

chmod(2), stat(2).

S-2





**NAME**

alarm — set a process alarm clock

**SYNOPSIS**

**unsigned alarm (sec)**  
**unsigned sec;**

**DESCRIPTION**

*Alarm* instructs the alarm clock of the calling process to send the signal **SIGALRM** to the calling process after the number of real time seconds specified by *sec* have elapsed; see *signal(2)*.

Alarm requests are not stacked; successive calls reset the alarm clock of the calling process.

If *sec* is 0, any previously made alarm request is canceled.

**RETURN VALUE**

*Alarm* returns the amount of time previously remaining in the alarm clock of the calling process.

**SEE ALSO**

pause(2), signal(2).

S-2

## BRK(2)

### NAME

*brk*, *sbrk* – change data segment space allocation

### SYNOPSIS

```
int brk (endds)
char *endds;
char *sbrk (incr)
int incr;
```

### DESCRIPTION

*Brk* and *sbrk* are used to change dynamically the amount of space allocated for the calling process's data segment(s); see *exec(2)*. The change is made by resetting the process's break value and allocating the appropriate amount of space. The *break value* is the address of the first data location beyond the end of allocated data. The amount of allocated space increases as the break value increases. The newly allocated space is set to zero.

*Brk* sets the break value to *endds* and changes the allocated space accordingly.

*Sbrk* adds *incr* bytes to the break value and changes the allocated space accordingly. *Incr* can be negative, in which case the amount of allocated space is decreased.

*Endds* and *incr* are rounded up to the next multiple of 512 in large and huge model programs.

*Brk* and *sbrk* will fail without making any change in the allocated space if one or more of the following are true:

- Such a change would result in more space being allocated than is allowed by a system-imposed maximum [see *ulimit(2)*]. [ENOMEM]

- Such a change would result in the segment selector of the break location being greater than or equal to the segment selector of any attached shared memory segment [see *shmop(2)*]. [ENOMEM]

- A large model process attempts to *brk* to an *endds* that has a segment selector which is greater than one more than the segment selector of the old break value. [ENOMEM]

- Such a change would result in the break value being in the stack or text areas of the process. [ENOMEM]

- Such a change would result in the break value being placed within an unallocated area between two currently allocated segments. [ENOMEM]

The following table summarizes the actions of *brk(2)*, and *sbrk(2)* in the different memory models (S = small, L = large, H = huge).

Operation	Model	Action
sbrk(0)	S	Returns current break value.
	L,H	Returns starting address of NEXT data segment.
sbrk(+incr)	S	Allocates incr bytes in current segment.
	L,H	Allocates incr bytes in next data segment (space from old break value* to end of old segment is not allocated).
	S,L,H	Returns the same value as sbrk(0).
sbrk(-incr)	S	Frees incr bytes in current segment.
	L,H	Frees incr bytes from as many segments as needed.
	S,L,H	Returns the same value as sbrk(0).
brk(endds) (current segment)	S,L,H	Sets break value to endds and allocates or frees memory to that point.
brk(endds) (previous segment)	L,H	Sets break value to endds and frees memory between old break value and endds. Endds must be an allocated location. Can free multiple segments.
brk(endds) (new segment)	L,H	Sets break value to endds in next segment.
	L	Can allocate up to one segment.
	H	Can allocate multiple segments.
	L,H	Space from old break value to end of old segment is not allocated.

\* "Old break value" is the break value previous to the execution of the current operation.

#### RETURN VALUE

Upon successful completion, *brk* returns a value of 0 and *sbrk* returns either the current break value (small model) or the starting address of next data segment (large and huge models). Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

#### CAVEATS

*Brk(2)* and *sbrk(2)* are not intended for general use: *Malloc(3C)* is the recommended way to obtain arbitrary amounts of memory.

Processes must not assume that the allocated address space is contiguous. When large and huge model processes perform any *sbrk* with a non-negative *incr* or a *brk* to a new segment, the area between the old segment's break location offset and the end of the old segment (offset 65535) is not accessible. Any reference to this area will cause a segmentation violation.

#### SEE ALSO

*exec(2)*, *shmop(2)*, *signal(2)*, *ulimit(2)*.

"Programming Procedures for UNIX System V/AT" in the Software Development System manual, Vol. I.

S-2

## CHDIR(2)

### NAME

`chdir` — change working directory

### SYNOPSIS

```
int chdir (path)
char *path;
```

### DESCRIPTION

*Path* points to the path name of a directory. *Chdir* causes the named directory to become the current working directory, the starting point for path searches for path names not beginning with */*.

*Chdir* will fail and the current working directory will be unchanged if one or more of the following are true:

- [ENOTDIR]      A component of the path name is not a directory.
- [ENOENT]      The named directory does not exist.
- [EACCES]      Search permission is denied for any component of the path name.
- [EFAULT]      *Path* points outside the allocated address space of the process.

### RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of `-1` is returned and *errno* is set to indicate the error.

### SEE ALSO

`chroot(2)`.

## NAME

chmod — change mode of file

## SYNOPSIS

```
int chmod (path, mode)
char *path;
int mode;
```

## DESCRIPTION

*Path* points to a path name naming a file. *Chmod* sets the access permission portion of the named file's mode according to the bit pattern contained in *mode*.

Access permission bits are interpreted as follows:

04000	Set user ID on execution.
02000	Set group ID on execution.
01000	Save text image after execution.
00400	Read by owner.
00200	Write by owner.
00100	Execute (search if a directory) by owner.
00070	Read, write, execute (search) by group.
00007	Read, write, execute (search) by others.

The effective user ID of the process must match the owner of the file or be super-user to change the mode of a file.

If the effective user ID of the process is not super-user, mode bit 01000 (save text image on execution) is cleared.

If the effective user ID of the process is not super-user and the effective group ID of the process does not match the group ID of the file, mode bit 02000 (set group ID on execution) is cleared.

If an executable file is prepared for sharing then mode bit 01000 prevents the system from abandoning the swap-space image of the program-text portion of the file when its last user terminates. Thus, when the next user of the file executes it, the text need not be read from the file system but can simply be swapped in, saving time.

*Chmod* will fail and the file mode will be unchanged if one or more of the following are true:

[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	The named file does not exist.
[EACCES]	Search permission is denied on a component of the path prefix.
[EPERM]	The effective user ID does not match the owner of the file and the effective user ID is not super-user.
[EROFS]	The named file resides on a read-only file system.
[EFAULT]	<i>Path</i> points outside the allocated address space of the process.

## RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of  $-1$  is returned and *errno* is set to indicate the error.

## SEE ALSO

chown(2), mknod(2).

## CHOWN(2)

### NAME

`chown` — change owner and group of a file

### SYNOPSIS

```
int chown (path, owner, group)
char *path;
int owner, group;
```

### DESCRIPTION

*Path* points to a path name naming a file. The owner ID and group ID of the named file are set to the numeric values contained in *owner* and *group* respectively.

Only processes with effective user ID equal to the file owner or super-user may change the ownership of a file.

If *chown* is invoked by other than the super-user, the set-user-ID and set-group-ID bits of the file mode, 04000 and 02000 respectively, will be cleared.

*Chown* will fail and the owner and group of the named file will remain unchanged if one or more of the following are true:

- [ENOTDIR]     A component of the path prefix is not a directory.
- [ENOENT]     The named file does not exist.
- [EACCES]     Search permission is denied on a component of the path prefix.
- [EPERM]      The effective user ID does not match the owner of the file and the effective user ID is not super-user.
- [EROFS]      The named file resides on a read-only file system.
- [EFAULT]     *Path* points outside the allocated address space of the process.

### RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

### SEE ALSO

`chmod(2)`,  
`chown(1)` in the Runtime System manual.

## NAME

chroot — change root directory

## SYNOPSIS

```
int chroot (path)
char *path;
```

## DESCRIPTION

*Path* points to a path name naming a directory. *Chroot* causes the named directory to become the root directory, the starting point for path searches for path names beginning with */*. The user's working directory is unaffected by the *chroot* system call.

The effective user ID of the process must be super-user to change the root directory.

The *..* entry in the root directory is interpreted to mean the root directory itself. Thus, *..* cannot be used to access files outside the subtree rooted at the root directory.

*Chroot* will fail and the root directory will remain unchanged if one or more of the following are true:

- [ENOTDIR] Any component of the path name is not a directory.
- [ENOENT] The named directory does not exist.
- [EPERM] The effective user ID is not super-user.
- [EFAULT] *Path* points outside the allocated address space of the process.

## RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

## SEE ALSO

chdir(2).

S-2

## CLOSE(2)

### NAME

close — close a file descriptor

### SYNOPSIS

```
int close (fdes)
int fdes;
```

### DESCRIPTION

*Fides* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, or *pipe* system call. *Close* closes the file descriptor indicated by *fides*.

*Close* will fail if *fides* is not a valid open file descriptor.

### RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of  $-1$  is returned and *errno* is set to indicate the error.

### SEE ALSO

*creat*(2), *dup*(2), *exec*(2), *fcntl*(2), *open*(2), *pipe*(2).



NAME

creat — create a new file or rewrite an existing one

SYNOPSIS

```
int creat (path, mode)
char *path;
int mode;
```

DESCRIPTION

*Creat* creates a new ordinary file or prepares to rewrite an existing file named by the path name pointed to by *path*.

If the file exists, the length is truncated to 0 and the mode and owner are unchanged. Otherwise, the file's owner ID is set to the effective user ID of the process, the group ID of the process is set to the effective group ID of the process, and the low-order 12 bits of the file mode are set to the value of *mode* modified as follows:

All bits set in the process's file mode creation mask are cleared. See *umask(2)*.

The "save text image after execution bit" of the mode is cleared. See *chmod(2)*.

Upon successful completion, the file descriptor is returned and the file is open for writing, even if the mode does not permit writing. The file pointer is set to the beginning of the file. The file descriptor is set to remain open across *exec* system calls. See *fcntl(2)*. No process may have more than 20 files open simultaneously. A new file may be created with a mode that forbids writing.

*Creat* will fail if one or more of the following are true:

- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] A component of the path prefix does not exist.
- [EACCES] Search permission is denied on a component of the path prefix.
- [ENOENT] The path name is null.
- [EACCES] The file does not exist and the directory in which the file is to be created does not permit writing.
- [EROFS] The named file resides or would reside on a read-only file system.
- [ETXTBSY] The file is a pure procedure (shared text) file that is being executed.
- [EACCES] The file exists and write permission is denied.
- [EISDIR] The named file is an existing directory.
- [EMFILE] Twenty (20) file descriptors are currently open.
- [EFAULT] *Path* points outside the allocated address space of the process.
- [ENFILE] The system file table is full.

RETURN VALUE

Upon successful completion, a non-negative integer, namely the file descriptor, is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO

*chmod(2)*, *close(2)*, *dup(2)*, *fcntl(2)*, *lseek(2)*, *open(2)*, *read(2)*, *umask(2)*, *write(2)*.

## DUP(2)

### NAME

`dup` — duplicate an open file descriptor

### SYNOPSIS

```
int dup (fildes)
int fildes;
```

### DESCRIPTION

*Fildes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, or *pipe* system call. *Dup* returns a new file descriptor having the following in common with the original:

Same open file (or pipe).

Same file pointer (i.e., both file descriptors share one file pointer).

Same access mode (read, write or read/write).

The new file descriptor is set to remain open across *exec* system calls. See *fcntl(2)*.

The file descriptor returned is the lowest one available.

*Dup* will fail if one or more of the following are true:

[EBADF] *Fildes* is not a valid open file descriptor.

[EMFILE] Twenty (20) file descriptors are currently open.

### RETURN VALUE

Upon successful completion a non-negative integer, namely the file descriptor, is returned. Otherwise, a value of `-1` is returned and *errno* is set to indicate the error.

### SEE ALSO

*creat(2)*, *close(2)*, *exec(2)*, *fcntl(2)*, *open(2)*, *pipe(2)*.

## NAME

exec, execv, execl, execve, execlp, execvp — execute a file

## SYNOPSIS

```
int execl (path, arg0, arg1, ..., argn, (char *)0)
char *path, *arg0, *arg1, ..., *argn;

int execv (path, argv)
char *path, *argv[];

int execl (path, arg0, arg1, ..., argn, (char *)0, envp)
char *path, *arg0, *arg1, ..., *argn, *envp[];

int execve (path, argv, envp)
char *path, *argv[], *envp[];

int execlp (file, arg0, arg1, ..., argn, (char *)0)
char *file, *arg0, *arg1, ..., *argn;

int execvp (file, argv)
char *file, *argv[];
```

## DESCRIPTION

*Exec* in all its forms transforms the calling process into a new process. The new process is constructed from an ordinary, executable file called the *new process file*. This file consists of a header [see *a.out*(4)], one or more text segments, and one or more data segments. The data segment may contain an initialized portion and an uninitialized portion (bss). There can be no return from a successful *exec* because the calling process is overlaid by the new process.

When a C program is executed, it is called as follows:

```
main (argc, argv, envp)
int argc;
char **argv, **envp;
```

where *argc* is the argument count and *argv* is an array of character pointers to the arguments themselves. As indicated, *argc* is conventionally at least one and the first member of the array points to a string containing the name of the file.

*Path* points to a path name that identifies the new process file.

*File* points to the new process file. The path prefix for this file is obtained by a search of the directories passed as the *environment* line "PATH =" [see *environ*(5)]. The environment is supplied by the shell [see *sh*(1)].

*Arg0*, *arg1*, ..., *argn* are pointers to null-terminated character strings. These strings constitute the argument list available to the new process. By convention, at least *arg0* must be present and point to a string that is the same as *path* (or its last component).

*Argv* is an array of character pointers to null-terminated strings. These strings constitute the argument list available to the new process. By convention, *argv* must have at least one member, and it must point to a string that is the same as *path* (or its last component). *Argv* is terminated by a null pointer.

*Envp* is an array of character pointers to null-terminated strings. These strings constitute the environment for the new process. *Envp* is terminated by a null pointer. For *execl* and *execv*, the C run-time start-off routine places a pointer to the environment of the calling process in the global cell:

```
extern char **environ;
```

and it is used to pass the environment of the calling process to the new process.

File descriptors open in the calling process remain open in the new process, except for those whose close-on-exec flag is set; see *fcntl*(2). For those file descriptors that remain open, the file pointer is unchanged.

## EXEC(2)

Signals set to terminate the calling process will be set to terminate the new process. Signals set to be ignored by the calling process will be set to be ignored by the new process. Signals set to be caught by the calling process will be set to terminate new process; see *signal(2)*.

If the set-user-ID mode bit of the new process file is set [see *chmod(2)*], *exec* sets the effective user ID of the new process to the owner ID of the new process file. Similarly, if the set-group-ID mode bit of the new process file is set, the effective group ID of the new process is set to the group ID of the new process file. The real user ID and real group ID of the new process remain the same as those of the calling process.

The shared memory segments attached to the calling process will not be attached to the new process [see *shmop(2)*].

Profiling is disabled for the new process; see *profil(2)*.

The new process also inherits the following attributes from the calling process:

- nice value [see *nice(2)*]
- process ID
- parent process ID
- process group ID
- semadj values [see *semop(2)*]
- tty group ID [see *exit(2)* and *signal(2)*]
- trace flag [see *ptrace(2)* request 0]
- time left until an alarm clock signal [see *alarm(2)*]
- current working directory
- root directory
- file mode creation mask [see *umask(2)*]
- file size limit [see *ulimit(2)*]
- utime*, *stime*, *cutime*, and *cstime* [see *times(2)*]

*Exec* will fail and return to the calling process if one or more of the following are true:

- [ENOENT] One or more components of the new process path name of the file do not exist.
- [ENOTDIR] A component of the new process path of the file prefix is not a directory.
- [EACCES] Search permission is denied for a directory listed in the new process file's path prefix.
- [EACCES] The new process file is not an ordinary file.
- [EACCES] The new process file mode denies execution permission.
- [ENOEXEC] The *exec* is not an *execlp* or *execvp*, and the new process file has the appropriate access permission but an invalid magic number in its header.
- [ETXTBSY] The new process file is a pure procedure (shared text) file that is currently open for writing by some process.
- [ENOMEM] The new process requires more memory than is allowed by the system-imposed maximum MAXMEM.
- [E2BIG] The number of bytes in the new process's argument list is greater than the system-imposed limit of 5120 bytes.
- [EFAULT] The new process file is not as long as indicated by the size values in its header.
- [EFAULT] *Path*, *argv*, or *envp* point to an illegal address.

**RETURN VALUE**

If *exec* returns to the calling process an error has occurred; the return value will be -1 and *errno* will be set to indicate the error.

**SEE ALSO**

alarm(2), exit(2), fork(2), nice(2), ptrace(2), semop(2), signal(2), times(2), ulimit(2), umask(2), a.out(4), environ(5), sh(1) in the Runtime System manual.

## EXIT(2)

### NAME

`exit`, `_exit` — terminate process

### SYNOPSIS

```
void exit (status)
int status;
void _exit (status)
int status;
```

### DESCRIPTION

*Exit* terminates the calling process with the following consequences:

All of the file descriptors open in the calling process are closed.

If the parent process of the calling process is executing a *wait*, it is notified of the calling process's termination and the low order eight bits (i.e., bits 0377) of *status* are made available to it; see *wait*(2).

If the parent process of the calling process is not executing a *wait*, the calling process is transformed into a zombie process. A *zombie process* is a process that only occupies a slot in the process table. It has no other space allocated either in user or kernel space. The process table slot that it occupies is partially overlaid with time accounting information (see `<sys/proc.h>`) to be used by *times*.

The parent process ID of all of the calling process's existing child processes and zombie processes is set to 1. This means the initialization process [see *intro*(2)] inherits each of these processes.

Each attached shared memory segment is detached and the value of `shm_nattach` in the data structure associated with its shared memory identifier is decremented by 1.

For each semaphore for which the calling process has set a `semadj` value [see *semop*(2)], that `semadj` value is added to the `semval` of the specified semaphore.

If the process has a process, text, or data lock, an *unlock* is performed [see *plock*(2)].

An accounting record is written on the accounting file if the system's accounting routine is enabled; see *acct*(2).

If the process ID, tty group ID, and process group ID of the calling process are equal, the *SIGHUP* signal is sent to each process that has a process group ID equal to that of the calling process.

The C function *exit* may cause cleanup actions before the process exits. The function *\_exit* circumvents all cleanup.

### SEE ALSO

*acct*(2), *intro*(2), *plock*(2), *semop*(2), *signal*(2), *wait*(2).

### WARNING

See *WARNING* in *signal*(2).

NAME  
fcntl - file control

SYNOPSIS  
#include <fcntl.h>  
int fcntl (fildes, cmd, arg)  
int fildes, cmd, arg;

DESCRIPTION  
*Fcntl* provides for control over open files. *Fildes* is an open file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, or *pipe* system call.

The *commands* available are:

F_DUPFD	Return a new file descriptor as follows: Lowest numbered available file descriptor greater than or equal to <i>arg</i> . Same open file (or pipe) as the original file. Same file pointer as the original file (i.e., both file descriptors share one file pointer). Same access mode (read, write, or read/write). Same file status flags (i.e., both file descriptors share the same file status flags). The close-on-exec flag associated with the new file descriptor is set to remain open across <i>exec(2)</i> system calls.
F_GETFD	Get the close-on-exec flag associated with the file descriptor <i>fildes</i> . If the low-order bit is 0, the file will remain open across <i>exec</i> ; otherwise, the file will be closed upon execution of <i>exec</i> .
F_SETFD	Set the close-on-exec flag associated with <i>fildes</i> to the low-order bit of <i>arg</i> (0 or 1 as above).
F_GETFL	Get <i>file</i> status flags.
F_SETFL	Set <i>file</i> status flags to <i>arg</i> . Only certain flags can be set; see <i>fcntl(5)</i> .
F_GETLK	Get the first lock which blocks the lock description given by the variable of type <i>struct flock</i> pointed to by <i>arg</i> . The information retrieved overwrites the information passed to <i>fcntl</i> in the <i>flock</i> structure. If no lock is found that would prevent this lock from being created, then the structure is passed back unchanged except for the lock type which will be set to F_UNLCK.
F_SETLK	Set or clear a file segment lock according to the variable of type <i>struct flock</i> pointed to by <i>arg</i> [see <i>fcntl(5)</i> ]. The <i>cmd</i> F_SETLK is used to establish read (F_RDLCK) and write (F_WRLCK) locks, as well as to remove either type of lock (F_UNLCK). If a read or write lock cannot be set, <i>fcntl</i> will return immediately with an error value of -1.
F_SETLKW	This <i>cmd</i> is the same as F_SETLK, except that if a read or write lock is blocked by other locks, the process will sleep until the segment is free to be locked.

## FCNTL(2)

A read lock prevents any process from write-locking the protected area. More than one read lock may exist for a given segment of a file at a given time. The file descriptor on which a read lock is being placed must have been opened with read access.

A write lock prevents any process from read-locking or write-locking the protected area. Only one write lock may exist for a given segment of a file at a given time. The file descriptor on which a write lock is being placed must have been opened with write access.

The structure *flock* describes the type (*l\_type*), starting offset (*l\_whence*), relative offset (*l\_start*), size (*l\_len*), and process id (*l\_pid*) of the segment of the file to be affected. The process id field is only used with the `F_GETLK` *cmd* to return the value for a blocking lock. Locks may start and extend beyond the current end of a file, but may not be negative relative to the beginning of the file. A lock may be set to always extend to the end of file by setting *l\_len* to zero (0). If such a lock also has *l\_start* set to zero (0), the whole file will be locked. Changing or unlocking a segment from the middle of a larger locked segment leaves two smaller segments for either end. Locking a segment that is already locked by the calling process causes the old lock type to be removed and the new lock type to take affect. All locks associated with a file for a given process are removed when a file descriptor for that file is closed by that process or the process holding that file descriptor terminates. Locks are not inherited by a child process in a `fork(2)` system call.

*Fcntl* will fail if one or more of the following are true:

- [EBADF] *Fildes* is not a valid open file descriptor.
- [EMFILE] *Cmd* is `F_DUPFD` and 20 file descriptors are currently open.
- [EINVAL] *Cmd* is `F_DUPFD` and *arg* is negative, greater than or equal to 20.
- [EINVAL] *Cmd* is `F_GETLK`, `F_SETLK`, or `F_SETLKW` and *arg* or the data it points to is not valid.
- [EACCESS] *Cmd* is `F_SETLK`, the type of lock (*l\_type*) is a read (`F_RDLCK`) or write (`F_WRLCK`) lock, and the segment of a file to be locked is already write-locked by another process, or the type is a write lock and the segment of a file to be locked is already read- or write-locked by another process.
- [EMFILE] *Cmd* is `F_SETLK` or `F_SETLKW`, the type of lock is a read or write lock and there are no more file-locking headers available (too many files have segments locked).
- [ENOSPC] *Cmd* is `F_SETLK` or `F_SETLKW`, the type of lock is a read or write lock and there are no more file-locking headers available (too many files have segments locked) or there are no more record locks available (too many file segments locked).
- [EDEADLK] *Cmd* is `F_SETLK`, when the lock is blocked by some lock from another process and sleeping (waiting) for that lock to become free; this causes a deadlock situation.

### SEE ALSO

`close(2)`, `exec(2)`, `open(2)`, `fcntl(5)`.



## RETURN VALUE

Upon successful completion, the value returned depends on *cmd* as follows:

F_DUPFD	A new file descriptor.
F_GETFD	Value of flag (only the low-order bit is defined).
F_SETFD	Value other than -1.
F_GETFL	Value of file flags.
F_SETFL	Value other than -1.
F_GETLK	Value other than -1.
F_SETLK	Value other than -1.
F_SETLKW	Value other than -1.

Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

## FORK(2)

### NAME

fork - create a new process

### SYNOPSIS

```
int fork ( )
```

### DESCRIPTION

*Fork* causes creation of a new process. The new process (child process) is an exact copy of the calling process (parent process). This means the child process inherits the following attributes from the parent process:

- environment
- close-on-exec flag [see *exec(2)*]
- signal-handling settings (i.e., **SIG\_DFL**, **SIG\_IGN**, function address)
- set-user-ID mode bit
- set-group-ID mode bit
- profiling on/off status
- nice value [see *nice(2)*]
- all attached shared memory segments [see *shmop(2)*]
- process group ID
- tty group ID [see *exit(2)* and *signal(2)*]
- trace flag [see *ptrace(2)* request 0]
- time left until an alarm clock signal [see *alarm(2)*]
- current working directory
- root directory
- file mode creation mask [see *umask(2)*]
- file size limit [see *ulimit(2)*]

The child process differs from the parent process in the following ways:

The child process has a unique process ID.

The child process has a different parent process ID (i.e., the process ID of the parent process).

The child process has its own copy of the parent's file descriptors. Each of the child's file descriptors shares a common file pointer with the corresponding file descriptor of the parent.

All *semadj* values are cleared [see *semop(2)*].

Process locks, text locks and data locks are not inherited by the child [see *plock(2)*].

The child process's *utime*, *stime*, *cutime*, and *cstime* are set to 0, the time left until an alarm clock signal is reset to 0.

*Fork* will fail and no child process will be created if one or more of the following are true:

[EAGAIN] The system-imposed limit on the total number of processes under execution would be exceeded.

[EAGAIN] The system-imposed limit on the total number of processes under execution by a single user would be exceeded.

### RETURN VALUE

Upon successful completion, *fork* returns a value of 0 to the child process and returns the process ID of the child process to the parent process. Otherwise, a value of -1 is returned to the parent process, no child process is created, and *errno* is set to indicate the error.

**NAME**

getpid, getpgrp, getppid — get process, process group, and parent process IDs

**SYNOPSIS**

int getpid ()

int getpgrp ()

int getppid ()

**DESCRIPTION**

*Getpid* returns the process ID of the calling process.

*Getpgrp* returns the process group ID of the calling process.

*Getppid* returns the parent process ID of the calling process.

**SEE ALSO**

exec(2), fork(2), intro(2), setpgrp(2), signal(2).

## GETUID(2)

### NAME

`getuid`, `geteuid`, `getgid`, `getegid` — get real user, effective user, real group, and effective group IDs

### SYNOPSIS

**unsigned short** `getuid` ()  
**unsigned short** `geteuid` ()  
**unsigned short** `getgid` ()  
**unsigned short** `getegid` ()

### DESCRIPTION

*Getuid* returns the real user ID of the calling process.

*Geteuid* returns the effective user ID of the calling process.

*Getgid* returns the real group ID of the calling process.

*Getegid* returns the effective group ID of the calling process.

### SEE ALSO

`intro(2)`, `setuid(2)`.

**NAME**

ioctl — control device

**SYNOPSIS**

```

ioctl (files, request, arg)
int files, request;
union ioctl_arg {
    int iarg; /*integer argument*/
    char *cparg; /*character pointer argument*/
} arg;

```

**DESCRIPTION**

*ioctl* performs a variety of functions on character special files (devices). The write-ups of various devices in Section 7 of the Runtime System manual discuss how *ioctl* applies to them.

*ioctl* will fail if one or more of the following are true:

- [EBADF]     *Files* is not a valid open file descriptor.
- [ENOTTY]    *Files* is not associated with a character special device.
- [EINVAL]    *Request* or *arg* is not valid. See Section 7 of the Runtime System manual.
- [EINTR]     A signal was caught during the *ioctl* system call.

**RETURN VALUE**

If an error has occurred, a value of -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

*termio(7)* in the Runtime System manual.

## KILL(2)

### NAME

kill — send a signal to a process or a group of processes

### SYNOPSIS

```
int kill (pid, sig)
```

```
int pid, sig;
```

### DESCRIPTION

*Kill* sends a signal to a process or a group of processes. The process or group of processes to which the signal is to be sent is specified by *pid*. The signal that is to be sent is specified by *sig* and is either one from the list given in *signal(2)*, or 0. If *sig* is 0 (the null signal), error checking is performed but no signal is actually sent. This can be used to check the validity of *pid*.

The real or effective user ID of the sending process must match the real or effective user ID of the receiving process, unless the effective user ID of the sending process is super-user.

The processes with a process ID of 0 and a process ID of 1 are special processes [see *intro(2)*] and will be referred to below as *proc0* and *proc1*, respectively.

If *pid* is greater than zero, *sig* will be sent to the process whose process ID is equal to *pid*. *Pid* may equal 1.

If *pid* is 0, *sig* will be sent to all processes excluding *proc0* and *proc1* whose process group ID is equal to the process group ID of the sender.

If *pid* is -1 and the effective user ID of the sender is not super-user, *sig* will be sent to all processes excluding *proc0* and *proc1* whose real user ID is equal to the effective user ID of the sender.

If *pid* is -1 and the effective user ID of the sender is super-user, *sig* will be sent to all processes excluding *proc0* and *proc1*.

If *pid* is negative but not -1, *sig* will be sent to all processes whose process group ID is equal to the absolute value of *pid*.

*Kill* will fail and no signal will be sent if one or more of the following are true:

[EINVAL] *Sig* is not a valid signal number.

[EINVAL] *Sig* is SIGKILL and *pid* is 1 (*proc1*).

[ESRCH] No process can be found corresponding to that specified by *pid*.

[EPERM] The user ID of the sending process is not super-user, and its real or effective user ID does not match the real or effective user ID of the receiving process.

### RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

### SEE ALSO

*getpid(2)*, *setpgid(2)*, *signal(2)*.

*kill(1)* in the Runtime System manual.

## NAME

link — link to a file

## SYNOPSIS

```
int link (path1, path2)
char *path1, *path2;
```

## DESCRIPTION

*Path1* points to a path name naming an existing file. *Path2* points to a path name naming the new directory entry to be created. *Link* creates a new link (directory entry) for the existing file.

*Link* will fail and no link will be created if one or more of the following are true:

- |           |                                                                                                                    |
|-----------|--------------------------------------------------------------------------------------------------------------------|
| [ENOTDIR] | A component of either path prefix is not a directory.                                                              |
| [ENOENT]  | A component of either path prefix does not exist.                                                                  |
| [EACCES]  | A component of either path prefix denies search permission.                                                        |
| [ENOENT]  | The file named by <i>path1</i> does not exist.                                                                     |
| [EEXIST]  | The link named by <i>path2</i> exists.                                                                             |
| [EPERM]   | The file named by <i>path1</i> is a directory and the effective user ID is not super-user.                         |
| [EXDEV]   | The link named by <i>path2</i> and the file named by <i>path1</i> are on different logical devices (file systems). |
| [ENOENT]  | <i>Path2</i> points to a null path name.                                                                           |
| [EACCES]  | The requested link requires writing in a directory with a mode that denies write permission.                       |
| [EROFS]   | The requested link requires writing in a directory on a read-only file system.                                     |
| [EFAULT]  | <i>Path</i> points outside the allocated address space of the process.                                             |
| [EMLINK]  | The maximum number of links to a file would be exceeded.                                                           |

## RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

## SEE ALSO

unlink(2).

## LSEEK(2)

### NAME

`lseek` — move read/write file pointer

### SYNOPSIS

```
long lseek (fildes, offset, whence)
int fildes;
long offset;
int whence;
```

### DESCRIPTION

*Fildes* is a file descriptor returned from a *creat*, *open*, *dup*, or *fcntl* system call. *Lseek* sets the file pointer associated with *fildes* as follows:

If *whence* is 0, the pointer is set to *offset* bytes.

If *whence* is 1, the pointer is set to its current location plus *offset*.

If *whence* is 2, the pointer is set to the size of the file plus *offset*.

Upon successful completion, the resulting pointer location, as measured in bytes from the beginning of the file, is returned.

*Lseek* will fail and the file pointer will remain unchanged if one or more of the following are true:

[EBADF] *Fildes* is not an open file descriptor.

[ESPIPE] *Fildes* is associated with a pipe or fifo.

[EINVAL and SIGSYS signal]  
*Whence* is not 0, 1, or 2.

[EINVAL] The resulting file pointer would be negative.

Some devices are incapable of seeking. The value of the file pointer associated with such a device is undefined.

### RETURN VALUE

Upon successful completion, a non-negative integer indicating the file pointer value is returned. Otherwise, a value of `-1` is returned and *errno* is set to indicate the error.

### SEE ALSO

`creat(2)`, `dup(2)`, `fcntl(2)`, `open(2)`.



## NAME

`mknod` — make a directory, or a special or ordinary file

## SYNOPSIS

```
int mknod (path, mode, dev)
char *path;
int mode, dev;
```

## DESCRIPTION

*Mknod* creates a new file named by the path name pointed to by *path*. The mode of the new file is initialized from *mode*. Where the value of *mode* is interpreted as follows:

```
0170000 file type; one of the following:
    0010000 fifo special
    0020000 character special
    0040000 directory
    0060000 block special
    0100000 or 0000000 ordinary file
0004000 set user ID on execution
0002000 set group ID on execution
0001000 save text image after execution
0000777 access permissions; constructed from the following
    0000400 read by owner
    0000200 write by owner
    0000100 execute (search on directory) by owner
    0000070 read, write, execute (search) by group
    0000007 read, write, execute (search) by others
```

The owner ID of the file is set to the effective user ID of the process. The group ID of the file is set to the effective group ID of the process.

Values of *mode* other than those above are undefined and should not be used. The low-order 9 bits of *mode* are modified by the process's file mode creation mask; all bits set in the process's file mode creation mask are cleared. See *umask*(2). If *mode* indicates a block or character special file, *dev* is a configuration-dependent specification of a character or block I/O device. If *mode* does not indicate a block special or character special device, *dev* is ignored.

*Mknod* may be invoked only by the super-user for file types other than FIFO special.

*Mknod* will fail and the new file will not be created if one or more of the following are true:

[EPERM]	The effective user ID of the process is not super-user.
[ENOTDIR]	A component of the path prefix is not a directory.
[ENOENT]	A component of the path prefix does not exist.
[EROFS]	The directory in which the file is to be created is located on a read-only file system.
[EEXIST]	The named file exists.
[EFAULT]	<i>Path</i> points outside the allocated address space of the process.

## **MKNOD(2)**

### **RETURN VALUE**

Upon successful completion a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

### **SEE ALSO**

`chmod(2)`, `exec(2)`, `umask(2)`, `fs(4)`,  
`mkdir(1)` in the Runtime System manual.

## NAME

mount - mount a file system

## SYNOPSIS

```
int mount (spec, dir, rwflag)
char *spec, *dir;
int rwflag;
```

## DESCRIPTION

*Mount* requests that a removable file system contained on the block special file identified by *spec* be mounted on the directory identified by *dir*. *Spec* and *dir* are pointers to path names.

Upon successful completion, references to the file *dir* will refer to the root directory on the mounted file system.

The low-order bit of *rwflag* is used to control write permission on the mounted file system; if 1, writing is forbidden, otherwise writing is permitted according to individual file accessibility.

*Mount* may be invoked only by the super-user.

*Mount* will fail if one or more of the following are true:

[EPERM]	The effective user ID is not super-user.
[ENOENT]	Any of the named files does not exist.
[ENOTDIR]	A component of a path prefix is not a directory.
[ENOTBLK]	<i>Spec</i> is not a block special device.
[ENXIO]	The device associated with <i>spec</i> does not exist.
[ENOTDIR]	<i>Dir</i> is not a directory.
[EFAULT]	<i>Spec</i> or <i>dir</i> points outside the allocated address space of the process.
[EBUSY]	<i>Dir</i> is currently mounted on, is someone's current working directory, or is otherwise busy.
[EBUSY]	The device associated with <i>spec</i> is currently mounted.
[EBUSY]	There are no more mount table entries.
[EROFS]	<i>Spec</i> is write-protected and <i>rwflag</i> requests write permission.
[ENOSPC]	The file system state in the super-block is not FsOKAY and <i>rwflag</i> requests write permission.
[EINVAL]	The file system magic is not FsMAGIC.

## RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

## SEE ALSO

umount(2), fs(4).

## MSGCTL(2)

### NAME

msgctl - message control operations

### SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/msg.h>

int msgctl (msqid, cmd, buf)
int msqid, cmd;
struct msqid_ds *buf;
```

### DESCRIPTION

*Msgctl* provides a variety of message control operations as specified by *cmd*. The following *cmds* are available:

**IPC\_STAT** Place the current value of each member of the data structure associated with *msqid* into the structure pointed to by *buf*. The contents of this structure are defined in *intro(2)*. {READ}

**IPC\_SET** Set the value of the following members of the data structure associated with *msqid* to the corresponding value found in the structure pointed to by *buf*:

- msg\_perm.uid
- msg\_perm.gid
- msg\_perm.mode /\* only low 9 bits \*/
- msg\_qbytes

This *cmd* can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of **msg\_perm.uid** in the data structure associated with *msqid*. Only super-user can raise the value of **msg\_qbytes**.

**IPC\_RMID** Remove the message queue identifier specified by *msqid* from the system and destroy the message queue and data structure associated with it. This *cmd* can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of **msg\_perm.uid** in the data structure associated with *msqid*.

*Msgctl* will fail if one or more of the following are true:

[EINVAL] *Msqid* is not a valid message queue identifier.

[EINVAL] *Cmd* is not a valid command.

[EACCES] *Cmd* is equal to **IPC\_STAT** and {READ} operation permission is denied to the calling process [see *intro(2)*].

[EPERM] *Cmd* is equal to **IPC\_RMID** or **IPC\_SET**. The effective user ID of the calling process is not equal to that of super-user and it is not equal to the value of **msg\_perm.uid** in the data structure associated with *msqid*.

[EPERM] *Cmd* is equal to **IPC\_SET**; an attempt is being made to increase to the value of **msg\_qbytes**, and the effective user ID of the calling process is not equal to that of super-user.

[EFAULT] *Buf* points to an illegal address.

## NAME

msgget — get message queue

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/msg.h>

int msgget (key, msgflg)
key_t key;
int msgflg;
```

## DESCRIPTION

*Msgget* returns the message queue identifier associated with *key*.

A message queue identifier and associated message queue and data structure [see *intro* (2)] are created for *key* if one of the following are true:

10 *Key* is equal to `IPC_PRIVATE`.

*Key* does not already have a message queue identifier associated with it, and (*msgflg* & `IPC_CREAT`) is “true”.

Upon creation, the data structure associated with the new message queue identifier is initialized as follows:

`Msg_perm.cuid`, `msg_perm.uid`, `msg_perm.cgid`, and `msg_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.

The low-order 9 bits of `msg_perm.mode` are set equal to the low-order 9 bits of *msgflg*.

`Msg_qnum`, `msg_lspid`, `msg_lrpid`, `msg_stime`, and `msg_rtime` are set equal to 0.

`Msg_ctime` is set equal to the current time.

`Msg_qbytes` is set equal to the system limit.

*Msgget* will fail if one or more of the following are true:

- |           |                                                                                                                                                                                |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [EACCESS] | A message queue identifier exists for <i>key</i> , but operation permission [see <i>intro</i> (2)] as specified by the low-order 9 bits of <i>msgflg</i> would not be granted. |
| [ENOENT]  | A message queue identifier does not exist for <i>key</i> and ( <i>msgflg</i> & <code>IPC_CREAT</code> ) is “false”.                                                            |
| [ENOSPC]  | A message queue identifier is to be created but the system-imposed limit on the maximum number of allowed message queue identifiers system wide would be exceeded.             |
| [EEXIST]  | A message queue identifier exists for <i>key</i> but ( ( <i>msgflg</i> & <code>IPC_CREAT</code> ) & ( <i>msgflg</i> & <code>IPC_EXCL</code> ) ) is “true”.                     |

## RETURN VALUE

Upon successful completion, a non-negative integer, namely a message queue identifier, is returned. Otherwise, a value of `-1` is returned and *errno* is set to indicate the error.

## SEE ALSO

*intro*(2), *msgctl*(2), *msgop*(2).

## MSGOP(2)

### NAME

msgop – message operations

### SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/msg.h>

int msgsnd (msqid, msgp, msgsz, msgflg)
int msqid;
struct msgbuf *msgp;
int msgsz, msgflg;

int msgrcv (msqid, msgp, msgsz, msgtyp, msgflg)
int msqid;
struct msgbuf *msgp;
int msgsz;
long msgtyp;
int msgflg;
```

### DESCRIPTION

Msgsnd is used to send a message to the queue associated with the message queue identifier specified by *msqid*. [WRITE] *Msgp* points to a structure containing the message. This structure is composed of the following members:

```
long    mtype;    /* message type */
char    mtext[]; /* message text */
```

*Mtype* is a positive integer that can be used by the receiving process for message selection (see *msgrcv* below). *Mtext* is any text of length *msgsz* bytes. *Msgsz* can range from 0 to a system-imposed maximum.

*Msgflg* specifies the action to be taken if one or more of the following are true:

The number of bytes already on the queue is equal to **msg\_qbytes** [see *intro(2)*].

The total number of messages on all queues system-wide is equal to the system-imposed limit.

These actions are as follows:

If (*msgflg* & **IPC\_NOWAIT**) is “true”, the message will not be sent and the calling process will return immediately.

If (*msgflg* & **IPC\_NOWAIT**) is “false”, the calling process will suspend execution until one of the following occurs:

The condition responsible for the suspension no longer exists, in which case the message is sent.

*Msqid* is removed from the system [see *msgctl(2)*]. When this occurs, *errno* is set equal to **EIDRM**, and a value of **-1** is returned.

The calling process receives a signal that is to be caught. In this case the message is not sent and the calling process resumes execution in the manner prescribed in *signal(2)*.

*Msgsnd* will fail and no message will be sent if one or more of the following are true:

[EINVAL]	<i>Msqid</i> is not a valid message queue identifier.
[EACCES]	Operation permission is denied to the calling process [see <i>intro(2)</i> ].

- [EINVAL] *Mtype* is less than 1.
- [EAGAIN] The message cannot be sent for one of the reasons cited above and (*msgflg* & *IPC\_NOWAIT*) is "true".
- [EINVAL] *Msgsz* is less than zero or greater than the system-imposed limit.
- [EFAULT] *Msgp* points to an illegal address.

Upon successful completion, the following actions are taken with respect to the data structure associated with *msgid* [see intro (2)].

*Msg\_qnum* is incremented by 1.

*Msg\_lspid* is set equal to the process ID of the calling process.

*Msg\_stime* is set equal to the current time.

*Msgrcv* reads a message from the queue associated with the message queue identifier specified by *msgid* and places it in the structure pointed to by *msgp* (READ) This structure is composed of the following members:

```
long    mtype;        /* message type */
char    mtext[];     /* message text */
```

*Mtype* is the received message's type as specified by the sending process. *Mtext* is the text of the message. *Msgsz* specifies the size in bytes of *mtext*. The received message is truncated to *msgsz* bytes if it is larger than *msgsz* and (*msgflg* & *MSG\_NOERROR*) is "true". The truncated part of the message is lost and no indication of the truncation is given to the calling process.

*Msgtyp* specifies the type of message requested as follows:

If *msgtyp* is equal to 0, the first message on the queue is received.

If *msgtyp* is greater than 0, the first message of type *msgtyp* is received.

If *msgtyp* is less than 0, the first message of the lowest type that is less than or equal to the absolute value of *msgtyp* is received.

*Msgflg* specifies the action to be taken if a message of the desired type is not on the queue. These are as follows:

If (*msgflg* & *IPC\_NOWAIT*) is "true", the calling process will return immediately with a return value of -1 and *errno* set to ENOMSG.

If (*msgflg* & *IPC\_NOWAIT*) is "false", the calling process will suspend execution until one of the following occurs:

A message of the desired type is placed on the queue.

*Msgid* is removed from the system. When this occurs, *errno* is set equal to EIDRM, and a value of -1 is returned.

The calling process receives a signal that is to be caught. In this case a message is not received and the calling process resumes execution in the manner prescribed in *signal(2)*.

*Msgrcv* will fail and no message will be received if one or more of the following are true:

- [EINVAL] *Msgid* is not a valid message queue identifier.
- [EACCES] Operation permission is denied to the calling process.
- [EINVAL] *Msgsz* is less than 0.
- [E2BIG] *Mtext* is greater than *msgsz* and (*msgflg* & *MSG\_NOERROR*) is "false".

## MSGOP(2)

[ENOMSG] The queue does not contain a message of the desired type and (*msgtyp* & `IPC_NOWAIT`) is "true".

[EFAULT] *Msgp* points to an illegal address.

Upon successful completion, the following actions are taken with respect to the data structure associated with *msqid* [see intro (2)].

**Msg\_qnum** is decremented by 1.

**Msg\_lrpid** is set equal to the process ID of the calling process.

**Msg\_rtime** is set equal to the current time.

### RETURN VALUES

If *msgsnd* or *msgrcv* return due to the receipt of a signal, a value of `-1` is returned to the calling process and *errno* is set to `EINTR`. If they return due to removal of *msqid* from the system, a value of `-1` is returned and *errno* is set to `EIDRM`.

Upon successful completion, the return value is as follows:

*Msgsnd* returns a value of 0.

*Msgrcv* returns a value equal to the number of bytes actually placed into *mtext*.

Otherwise, a value of `-1` is returned and *errno* is set to indicate the error.

### SEE ALSO

`intro(2)`, `msgctl(2)`, `msgget(2)`, `signal(2)`.



**NAME**

*nice* - change priority of a process

**SYNOPSIS**

```
int nice (incr)
int incr;
```

**DESCRIPTION**

*Nice* adds the value of *incr* to the nice value of the calling process. A process's *nice value* is a positive number for which a more positive value results in lower CPU priority.

A maximum nice value of 39 and a minimum nice value of 0 are imposed by the system. Requests for values above or below these limits result in the nice value being set to the corresponding limit.

[EPERM] *Nice* will fail and not change the nice value if *incr* is negative or greater than 40 and the effective user ID of the calling process is not super-user.

**RETURN VALUE**

Upon successful completion, *nice* returns the new nice value minus 20. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

*exec(2)*  
*nice(1)* in the Runtime System manual.

# OPEN(2)

## NAME

open - open for reading or writing

## SYNOPSIS

```
#include <fcntl.h>
int open (path, oflag [ , mode ] )
char *path;
int oflag, mode;
```

## DESCRIPTION

*Path* points to a path name naming a file. *Open* opens a file descriptor for the named file and sets the file status flags according to the value of *oflag*. *Oflag* values are constructed by or-ing flags from the following list (only one of the first three flags below may be used):

**O\_RDONLY** Open for reading only.

**O\_WRONLY** Open for writing only.

**O\_RDWR** Open for reading and writing.

**O\_NDELAY** This flag may affect subsequent reads and writes. See *read(2)* and *write(2)*.

When opening a FIFO with **O\_RDONLY** or **O\_WRONLY** set:

If **O\_NDELAY** is set:

An *open* for reading-only will return without delay. An *open* for writing-only will return an error if no process currently has the file open for reading.

If **O\_NDELAY** is clear:

An *open* for reading-only will block until a process opens the file for writing. An *open* for writing-only will block until a process opens the file for reading.

When opening a file associated with a communication line:

If **O\_NDELAY** is set:

The open will return without waiting for carrier.

If **O\_NDELAY** is clear:

The open will block until carrier is present.

**O\_APPEND** If set, the file pointer will be set to the end of the file prior to each write.

**O\_SYNC** When opening a regular file, this flag affects subsequent writes. If set, each *write(2)* will wait for both the file data and file status to be physically updated.

**O\_CREAT** If the file exists, this flag has no effect. Otherwise, the owner ID of the file is set to the effective user ID of the process, the group ID of the file is set to the effective group ID of the process, and the low-order 12 bits of the file mode are set to the value of *mode* modified as follows [see *creat(2)*]:

All bits set in the file mode creation mask of the process are cleared. See *umask(2)*.

The "save text image after execution bit" of the mode is cleared. See *chmod(2)*.

- O\_TRUNC** If the file exists, its length is truncated to 0 and the mode and owner are unchanged.
- O\_EXCL** If **O\_EXCL** and **O\_CREAT** are set, *open* will fail if the file exists.

The file pointer used to mark the current position within the file is set to the beginning of the file.

The new file descriptor is set to remain open across *exec* system calls. See *fcntl(2)*.

The named file is opened unless one or more of the following are true:

- [ENOTDIR] A component of the path prefix is not a directory.
- [ENOENT] **O\_CREAT** is not set and the named file does not exist.
- [EACCES] A component of the path prefix denies search permission.
- [EACCES] *Oflag* permission is denied for the named file.
- [EISDIR] The named file is a directory and *oflag* is write or read/write.
- [EROFS] The named file resides on a read-only file system and *oflag* is write or read/write.
- [EMFILE] Twenty (20) file descriptors are currently open.
- [ENXIO] The named file is a character special or block special file, and the device associated with this special file does not exist.
- [ETXTBSY] The file is a pure procedure (shared text) file that is being executed, and *oflag* is write or read/write.
- [EFAULT] *Path* points outside the allocated address space of the process.
- [EEXIST] **O\_CREAT** and **O\_EXCL** are set, and the named file exists.
- [ENXIO] **O\_NDELAY** is set, the named file is a FIFO, **O\_WRONLY** is set, and no process has the file open for reading.
- [EINTR] A signal was caught during the *open* system call.
- [ENFILE] The system file table is full.

#### RETURN VALUE

Upon successful completion, the file descriptor is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

#### SEE ALSO

*chmod(2)*, *close(2)*, *creat(2)*, *dup(2)*, *fcntl(2)*, *lseek(2)*, *read(2)*, *umask(2)*, *write(2)*.

## PAUSE(2)

### NAME

pause - suspend process until signal

### SYNOPSIS

**pause ( )**

### DESCRIPTION

*Pause* suspends the calling process until it receives a signal. The signal must be one that is not currently set to be ignored by the calling process.

If the signal causes termination of the calling process, *pause* will not return.

If the signal is *caught* by the calling process and control is returned from the signal-catching function [see *signal(2)*], the calling process resumes execution from the point of suspension with a return value of -1 from *pause* and *errno* set to EINTR.

### SEE ALSO

alarm(2), kill(2), signal(2), wait(2).

**NAME**

pipe — create an interprocess channel

**SYNOPSIS**

```
int pipe (fildes)
int fildes[2];
```

**DESCRIPTION**

*Pipe* creates an I/O mechanism called a pipe and returns two file descriptors, *fildes*[0] and *fildes*[1]. *Fildes*[0] is opened for reading and *fildes*[1] is opened for writing.

Up to 5120 bytes of data are buffered by the pipe before the writing process is blocked. A read only file descriptor *fildes*[0] accesses the data written to *fildes*[1] on a first-in-first-out (FIFO) basis.

[EMFILE] *Pipe* will fail if 19 or more file descriptors are currently open.

[ENFILE] The system file table is full.

**RETURN VALUE**

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

*read*(2), *write*(2),  
*sh*(1) in the Runtime System manual.

## PLOCK(2)

### NAME

`plock` — lock process, text, or data in memory

### SYNOPSIS

```
#include <sys/lock.h>
```

```
int plock (op)
```

```
int op;
```

### DESCRIPTION

`Plock` allows the calling process to lock its text segment (text lock), its data segment (data lock), or both its text and data segments (process lock) into memory. Locked segments are immune to all routine swapping. `Plock` also allows these segments to be unlocked. The effective user ID of the calling process must be super-user to use this call. `Op` specifies the following:

**PROCLOCK** — lock text and data segments into memory (process lock)

**TXTLCK** — lock text segment into memory (text lock)

**DATLOCK** — lock data segment into memory (data lock)

**UNLOCK** — remove locks

`Plock` will fail and not perform the requested operation if one or more of the following are true:

[EPERM] The effective user ID of the calling process is not super-user.

[EINVAL] `Op` is equal to **PROCLOCK** and a process lock, a text lock, or a data lock already exists on the calling process.

[EINVAL] `Op` is equal to **TXTLCK** and a text lock, or a process lock already exists on the calling process.

[EINVAL] `Op` is equal to **DATLOCK** and a data lock, or a process lock already exists on the calling process.

[EINVAL] `Op` is equal to **UNLOCK** and no type of lock exists on the calling process.

### RETURN VALUE

Upon successful completion, a value of 0 is returned to the calling process. Otherwise, a value of -1 is returned and `errno` is set to indicate the error.

### SEE ALSO

`exec(2)`, `exit(2)`, `fork(2)`.

**NAME**

profil — execution time profile

**SYNOPSIS**

```
void profil (buff, bufsiz, offset, scale)
short *buff;
int bufsiz, scale;
void (*offset)();
```

**DESCRIPTION**

*Buff* points to an area of core whose length (in bytes) is given by *bufsiz*. After this call, the user's program counter (*pc*) is examined each clock tick (60th second); *offset* is subtracted from it, and the result multiplied by *scale*. If the resulting number corresponds to a word inside *buff*, that word is incremented.

The scale is interpreted as an unsigned, fixed-point fraction with binary point at the left: 0177777 (octal) gives a 1-1 mapping of *pc*'s to words in *buff*; 077777 (octal) maps each pair of instruction words together. 02(octal) maps all instructions onto the beginning of *buff* (producing a non-interrupting core clock).

Profiling is turned off by giving a *scale* of 0 or 1. It is rendered ineffective by giving a *bufsiz* of 0. Profiling is turned off when an *exec* is executed, but remains on in child and parent both after a *fork*. Profiling will be turned off if an update in *buff* would cause a memory fault.

**RETURN VALUE**

Not defined.

**SEE ALSO**

monitor(3C).  
 prof(1) in the Runtime System manual.

## PTRACE(2)

### NAME

ptrace — process trace

### SYNOPSIS

```
int ptrace (request, pid, addr, data);
int request, pid, data;
char *addr;
```

### DESCRIPTION

*Ptrace* provides a means by which a parent process may control the execution of a child process. Its primary use is for the implementation of breakpoint debugging; see *sdb*(1). The child process behaves normally until it encounters a signal [see *signal*(2) for the list], at which time it enters a stopped state and its parent is notified via *wait*(2). When the child is in the stopped state, its parent can examine and modify its “core image” using *ptrace*. Also, the parent can cause the child either to terminate or continue, with the possibility of ignoring the signal that caused it to stop.

The *request* argument determines the precise action to be taken by *ptrace* and is one of the following:

- 0 This request must be issued by the child process if it is to be traced by its parent. It turns on the child's trace flag that stipulates that the child should be left in a stopped state upon receipt of a signal rather than the state specified by *func*; see *signal*(2). The *pid*, *addr*, and *data* arguments are ignored, and a return value is not defined for this request. Peculiar results will ensue if the parent does not expect to trace the child.

The remainder of the requests can only be used by the parent process. For each, *pid* is the process ID of the child. The child must be in a stopped state before these requests are made.

- 1, 2 With these requests, the word at location *addr* in the address space of the child is returned to the parent process. If I and D space are separated (as on iAPX 286s and PDP-11s), request 1 returns a word from I space, and request 2 returns a word from D space. If I and D space are not separated (as on the 3B20 computer and VAX-11/780), either request 1 or request 2 may be used with equal results. The *data* argument is ignored. These two requests will fail if *addr* is not the start address of a word, in which case a value of -1 is returned to the parent process and the parent's *errno* is set to EIO.
- 3 With this request, the word at location *addr* in the child's USER area in the system's address space (see `<sys/user.h>`) is returned to the parent process. Addresses in this area range from 0 to 1024 on the PDP-11s and 0 to 2048 on the 3B20 computer, the iAPX 286 and VAX. The *data* argument is ignored. This request will fail if *addr* is not the start address of a word or is outside the USER area, in which case a value of -1 is returned to the parent process and the parent's *errno* is set to EIO.
- 4, 5 With these requests, the value given by the *data* argument is written into the address space of the child at location *addr*. If I and D space are separated (as on the iAPX 286 and PDP-11s), request 4 writes a word into I space, and request 5 writes a word into D space. If I and D space are not separated (as on the 3B20 computer and VAX), either request 4 or request 5 may be used with equal results. Upon successful completion, the value written into the address space of the child is returned to the parent.



These two requests will fail if *addr* is a location in a pure procedure space and another process is executing in that space, or *addr* is not the start address of a word. Upon failure a value of -1 is returned to the parent process and the parent's *errno* is set to EIO.

- 6 With this request, a few entries in the child's USER area can be written. *Data* gives the value that is to be written and *addr* is the location of the entry. The few entries that can be written are:

the general registers (i.e., registers 0-11 on the 3B20 computer, registers 0-7 on PDP-11s, all registers except SS on the iAPX 286 and registers 0-15 on the VAX)

the condition codes of the Processor Status Word on the 3B20 computer

the floating point status register and six floating point registers on PDP-11s

certain bits of the Processor Status Word on PDP-11s (i.e., bits 0-4, and 8-11)

certain bits of the flags register on the iAPX 286 (i.e., bits 0-7, and 10-11)

certain bits of the Processor Status Longword on the VAX (i.e., bits 0-7, 16-20, and 30-31).

- 7 This request causes the child to resume execution. If the *data* argument is 0, all pending signals including the one that caused the child to stop are canceled before it resumes execution. If the *data* argument is a valid signal number, the child resumes execution as if it had incurred that signal, and any other pending signals are canceled. The *addr* argument must be equal to 1 for this request. Upon successful completion, the value of *data* is returned to the parent. This request will fail if *data* is not 0 or a valid signal number, in which case a value of -1 is returned to the parent process and the parent's *errno* is set to EIO.
- 8 This request causes the child to terminate with the same consequences as *exit(2)*.
- 9 This request sets the trace bit in the Processor Status Word of the child (i.e., bit 4 on PDP-11s, bit 8 on the iAPX 286, and bit 30 on the VAX) and then executes the same steps as listed above for request 7. The trace bit causes an interrupt upon completion of one machine instruction. This effectively allows single stepping of the child. On the 3B20 computer there is no trace bit and this request returns an error.

Note: the trace bit remains set after an interrupt on PDP-11s but is turned off after an interrupt on the iAPX 286, or the VAX.

To forestall possible fraud, *ptrace* inhibits the set-user-id facility on subsequent *exec(2)* calls. If a traced process calls *exec*, it will stop before executing the first instruction of the new image showing signal SIGTRAP.

#### GENERAL ERRORS

*Ptrace* will in general fail if one or more of the following are true:

- [EIO] *Request* is an illegal number.
- [ESRCH] *Pid* identifies a child that does not exist or has not executed a *ptrace* with request 0.

## **PTRACE(2)**

### **CAVEAT**

It is not generally possible for a small model process to trace a large or huge model process.

### **SEE ALSO**

`exec(2)`, `signal(2)`, `wait(2)`.  
`sdb(1)` in the Runtime System manual.

## NAME

read — read from file

## SYNOPSIS

```
int read (fildes, buf, nbyte)
int fildes;
char *buf;
unsigned nbyte;
```

## DESCRIPTION

*Fildes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, or *pipe* system call.

*Read* attempts to read *nbyte* bytes from the file associated with *fildes* into the buffer pointed to by *buf*.

On devices capable of seeking, the *read* starts at a position in the file given by the file pointer associated with *fildes*. Upon return from *read*, the file pointer is incremented by the number of bytes actually read.

Devices that are incapable of seeking always read from the current position. The value of a file pointer associated with such a file is undefined.

Upon successful completion, *read* returns the number of bytes actually read and placed in the buffer; this number may be less than *nbyte* if the file is associated with a communication line [see *ioctl(2)* and *termio(7)*], or if the number of bytes left in the file is less than *nbyte* bytes. A value of 0 is returned when an end-of-file has been reached.

When attempting to read from an empty pipe (or FIFO):

If `O_NDELAY` is set, the read will return a 0.

If `O_NDELAY` is clear, the read will block until data is written to the file or the file is no longer open for writing.

When attempting to read a file associated with a tty that has no data currently available:

If `O_NDELAY` is set, the read will return a 0.

If `O_NDELAY` is clear, the read will block until data becomes available.

*Read* will fail if one or more of the following are true:

- [EBADF] *Fildes* is not a valid file descriptor open for reading.
- [EFAULT] *Buf* points outside the allocated address space.
- [EINTR] A signal was caught during the *read* system call.

## RETURN VALUE

Upon successful completion a non-negative integer is returned indicating the number of bytes actually read. Otherwise, a `-1` is returned and *errno* is set to indicate the error.

## SEE ALSO

*creat(2)*, *dup(2)*, *fcntl(2)*, *ioctl(2)*, *open(2)*, *pipe(2)*, *termio(7)* in the Runtime System manual.

## SEMCTL(2)

### NAME

semctl — semaphore control operations

### SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>

int semctl (semid, semnum, cmd, arg)
int semid, cmd;
int semnum;
union semun {
    int val;
    struct semid_ds *buf;
    ushort *array;
} arg;
```

### DESCRIPTION

*Semctl* provides a variety of semaphore control operations as specified by *cmd*.

The following *cmds* are executed with respect to the semaphore specified by *semid* and *semnum*:

- |         |                                                                                                                                                                                                         |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GETVAL  | Return the value of <i>semval</i> [see <i>intro(2)</i> ]. {READ}                                                                                                                                        |
| SETVAL  | Set the value of <i>semval</i> to <i>arg.val</i> . {ALTER} When this <i>cmd</i> is successfully executed, the <i>semadj</i> value corresponding to the specified semaphore in all processes is cleared. |
| GETPID  | Return the value of <i>sempid</i> . {READ}                                                                                                                                                              |
| GETNCNT | Return the value of <i>semncnt</i> . {READ}                                                                                                                                                             |
| GETZCNT | Return the value of <i>semzcnt</i> . {READ}                                                                                                                                                             |

The following *cmds* return and set, respectively, every *semval* in the set of semaphores.

- |        |                                                                                                                                                                                                                                    |
|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GETALL | Place <i>semvals</i> into array pointed to by <i>arg.array</i> . {READ}                                                                                                                                                            |
| SETALL | Set <i>semvals</i> according to the array pointed to by <i>arg.array</i> . {ALTER} When this <i>cmd</i> is successfully executed, the <i>semadj</i> values corresponding to each specified semaphore in all processes are cleared. |

The following *cmds* are also available:

- |          |                                                                                                                                                                                                                                                                        |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPC_STAT | Place the current value of each member of the data structure associated with <i>semid</i> into the structure pointed to by <i>arg.buf</i> . The contents of this structure are defined in <i>intro(2)</i> . {READ}                                                     |
| IPC_SET  | Set the value of the following members of the data structure associated with <i>semid</i> to the corresponding value found in the structure pointed to by <i>arg.buf</i> :<br><b>sem_perm.uid</b><br><b>sem_perm.gid</b><br><b>sem_perm.mode</b> /* only low 9 bits */ |

This *cmd* can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of *sem\_perm.uid* in the data structure associated with *semid*.

**IPC\_RMID** Remove the semaphore identifier specified by *semid* from the system and destroy the set of semaphores and data structure associated with it. This *cmd* can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of *sem\_perm.uid* in the data structure associated with *semid*.

*Semctl* will fail if one or more of the following are true:

- [EINVAL] *Semid* is not a valid semaphore identifier.
- [EINVAL] *Semnum* is less than zero or greater than *sem\_nsems*.
- [EINVAL] *Cmd* is not a valid command.
- [EACCES] Operation permission is denied to the calling process [see *intro(2)*].
- [ERANGE] *Cmd* is SETVAL or SETALL and the value to which *semval* is to be set is greater than the system imposed maximum.
- [EPERM] *Cmd* is equal to IPC\_RMID or IPC\_SET and the effective user ID of the calling process is not equal to that of super-user and it is not equal to the value of *sem\_perm.uid* in the data structure associated with *semid*.
- [EFAULT] *Arg.buf* points to an illegal address.

#### RETURN VALUE

Upon successful completion, the value returned depends on *cmd* as follows:

- GETVAL The value of *semval*.
- GETPID The value of *sempid*.
- GETNCNT The value of *semncnt*.
- GETZCNT The value of *semzcnt*.
- All others A value of 0.

Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

#### SEE ALSO

*intro(2)*, *semget(2)*, *semop(2)*.

## SEMGET(2)

### NAME

semget — get set of semaphores

### SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>

int semget (key, nsems, semflg)
key_t key;
int nsems, semflg;
```

### DESCRIPTION

*Semget* returns the semaphore identifier associated with *key*.

A semaphore identifier and associated data structure and set containing *nsems* semaphores [see *intro(2)*] are created for *key* if one of the following are true:

*Key* is equal to `IPC_PRIVATE`.

*Key* does not already have a semaphore identifier associated with it, and (*semflg* & `IPC_CREAT`) is “true”.

Upon creation, the data structure associated with the new semaphore identifier is initialized as follows:

`Sem_perm.cuid`, `sem_perm.uid`, `sem_perm.cgid`, and `sem_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.

The low-order 9 bits of `sem_perm.mode` are set equal to the low-order 9 bits of *semflg*.

`Sem_nsems` is set equal to the value of *nsems*.

`Sem_otime` is set equal to 0 and `sem_ctime` is set equal to the current time.

*Semget* will fail if one or more of the following are true:

- |          |                                                                                                                                                                            |
|----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [EINVAL] | <i>Nsems</i> is either less than or equal to zero or greater than the system-imposed limit.                                                                                |
| [EACCES] | A semaphore identifier exists for <i>key</i> , but operation permission [see <i>intro(2)</i> ] as specified by the low-order 9 bits of <i>semflg</i> would not be granted. |
| [EINVAL] | A semaphore identifier exists for <i>key</i> , but the number of semaphores in the set associated with it is less than <i>nsems</i> and <i>nsems</i> is not equal to zero. |
| [ENOENT] | A semaphore identifier does not exist for <i>key</i> and ( <i>semflg</i> & <code>IPC_CREAT</code> ) is “false”.                                                            |
| [ENOSPC] | A semaphore identifier is to be created but the system-imposed limit on the maximum number of allowed semaphore identifiers system wide would be exceeded.                 |
| [ENOSPC] | A semaphore identifier is to be created but the system-imposed limit on the maximum number of allowed semaphores system wide would be exceeded.                            |
| [EEXIST] | A semaphore identifier exists for <i>key</i> but ( ( <i>semflg</i> & <code>IPC_CREAT</code> ) and ( <i>semflg</i> & <code>IPC_EXCL</code> ) ) is “true”.                   |

**RETURN VALUE**

Upon successful completion, a non-negative integer, namely a semaphore identifier, is returned. Otherwise, a value of  $-1$  is returned and *errno* is set to indicate the error.

**SEE ALSO**

intro(2), semctl(2), semop(2).

## SEMOP(2)

### NAME

semop - semaphore operations

### SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/sem.h>

int semop (semid, sops, nsops)
int semid;
struct sembuf **sops;
int nsops;
```

### DESCRIPTION

*Semop* is used to automatically perform an array of semaphore operations on the set of semaphores associated with the semaphore identifier specified by *semid*. *Sops* is a pointer to the array of semaphore-operation structures. *Nsops* is the number of such structures in the array. The contents of each structure includes the following members:

```
short   sem_num; /* semaphore number */
short   sem_op; /* semaphore operation */
short   sem_flg; /* operation flags */
```

Each semaphore operation specified by *sem\_op* is performed on the corresponding semaphore specified by *semid* and *sem\_num*.

*Sem\_op* specifies one of three semaphore operations as follows:

If *sem\_op* is a negative integer, one of the following will occur:  
{ALTER}

If *semval* [see *intro(2)*] is greater than or equal to the absolute value of *sem\_op*, the absolute value of *sem\_op* is subtracted from *semval*. Also, if (*sem\_flg* & SEM\_UNDO) is "true", the absolute value of *sem\_op* is added to the calling process's *semadj* value [see *exit(2)*] for the specified semaphore.

If *semval* is less than the absolute value of *sem\_op* and (*sem\_flg* & IPC\_NOWAIT) is "true", *semop* will return immediately.

If *semval* is less than the absolute value of *sem\_op* and (*sem\_flg* & IPC\_NOWAIT) is "false", *semop* will increment the *semncnt* associated with the specified semaphore and suspend execution of the calling process until one of the following conditions occur.

*Semval* becomes greater than or equal to the absolute value of *sem\_op*. When this occurs, the value of *semncnt* associated with the specified semaphore is decremented, the absolute value of *sem\_op* is subtracted from *semval* and, if (*sem\_flg* & SEM\_UNDO) is "true", the absolute value of *sem\_op* is added to the calling process's *semadj* value for the specified semaphore.

The *semid* for which the calling process is awaiting action is removed from the system [see *semctl(2)*]. When this occurs, *errno* is set equal to EIDRM, and a value of -1 is returned.



The calling process receives a signal that is to be caught. When this occurs, the value of `semncnt` associated with the specified semaphore is decremented, and the calling process resumes execution in the manner prescribed in *signal(2)*.

If `sem_op` is a positive integer, the value of `sem_op` is added to `semval` and, if `(sem_flg & SEM_UNDO)` is "true", the value of `sem_op` is subtracted from the calling process's `semadj` value for the specified semaphore. [ALTER]

If `sem_op` is zero, one of the following will occur: [READ]

If `semval` is zero, `semop` will return immediately.

If `semval` is not equal to zero and `(sem_flg & IPC_NOWAIT)` is "true", `semop` will return immediately.

If `semval` is not equal to zero and `(sem_flg & IPC_NOWAIT)` is "false", `semop` will increment the `semzcnt` associated with the specified semaphore and suspend execution of the calling process until one of the following occurs:

`Semval` becomes zero, at which time the value of `semzcnt` associated with the specified semaphore is decremented.

The `semid` for which the calling process is awaiting action is removed from the system. When this occurs, `errno` is set equal to `EIDRM`, and a value of `-1` is returned.

The calling process receives a signal that is to be caught. When this occurs, the value of `semzcnt` associated with the specified semaphore is decremented, and the calling process resumes execution in the manner prescribed in *signal(2)*.

`Semop` will fail if one or more of the following are true for any of the semaphore operations specified by `sops`:

- [EINVAL] `Semid` is not a valid semaphore identifier.
- [EFBIG] `Sem_num` is less than zero or greater than or equal to the number of semaphores in the set associated with `semid`.
- [E2BIG] `Nsops` is greater than the system-imposed maximum.
- [EACCES] Operation permission is denied to the calling process [see *intro(2)*].
- [EAGAIN] The operation would result in suspension of the calling process but `(sem_flg & IPC_NOWAIT)` is "true".
- [ENOSPC] The limit on the number of individual processes requesting an `SEM_UNDO` would be exceeded.
- [EINVAL] The number of individual semaphores for which the calling process requests a `SEM_UNDO` would exceed the limit.
- [ERANGE] An operation would cause a `semval` to overflow the system-imposed limit.
- [ERANGE] An operation would cause a `semadj` value to overflow the system-imposed limit.
- [EFAULT] `Sops` points to an illegal address.

Upon successful completion, the value of `sempid` for each semaphore specified in the array pointed to by `sops` is set equal to the process ID of the calling process.

## SEMOP(2)

### RETURN VALUE

If *semop* returns due to the receipt of a signal, a value of  $-1$  is returned to the calling process and *errno* is set to EINTR. If it returns due to the removal of a *semid* from the system, a value of  $-1$  is returned and *errno* is set to EIDRM.

Upon successful completion, the value of *semval* at the time of the call for the last operation in the array pointed to by *sops* is returned. Otherwise, a value of  $-1$  is returned and *errno* is set to indicate the error.

### SEE ALSO

*exec(2)*, *exit(2)*, *fork(2)*, *intro(2)*, *semctl(2)*, *semget(2)*.

**NAME**

setpgrp — set process group ID

**SYNOPSIS**

**int** setpgrp (**)**

**DESCRIPTION**

*Setpgrp* sets the process group ID of the calling process to the process ID of the calling process and returns the new process group ID.

**RETURN VALUE**

*Setpgrp* returns the value of the new process group ID.

**SEE ALSO**

exec(2), fork(2), getpid(2), intro(2), kill(2), signal(2).

## SETUID(2)

### NAME

setuid, setgid — set user and group IDs

### SYNOPSIS

```
int setuid (uid)
int uid;
int setgid (gid)
int gid;
```

### DESCRIPTION

*Setuid (setgid)* is used to set the real user (group) ID and effective user (group) ID of the calling process.

If the effective user ID of the calling process is super-user, the real user (group) ID and effective user (group) ID are set to *uid (gid)*.

If the effective user ID of the calling process is not super-user, but its real user (group) ID is equal to *uid (gid)*, the effective user (group) ID is set to *uid (gid)*.

If the effective user ID of the calling process is not super-user, but the saved set-user (group) ID from *exec(2)* is equal to *uid (gid)*, the effective user (group) ID is set to *uid (gid)*.

*Setuid (setgid)* will fail if the real user (group) ID of the calling process is not equal to *uid (gid)* and its effective user ID is not super-user. [EPERM]

The *uid* is out of range. [EINVAL]

### RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

### SEE ALSO

getuid(2), intro(2).

## NAME

shmctl — shared memory control operations

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>

int shmctl (shmid, cmd, buf)
int shmid, cmd;
struct shmctl *buf;
```

## DESCRIPTION

*Shmctl* provides a variety of shared memory control operations as specified by *cmd*. The following *cmds* are available:

- |          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IPC_STAT | Place the current value of each member of the data structure associated with <i>shmid</i> into the structure pointed to by <i>buf</i> . The contents of this structure are defined in <i>intro</i> (2). (READ)                                                                                                                                                                                                                                                          |
| IPC_SET  | Set the value of the following members of the data structure associated with <i>shmid</i> to the corresponding value found in the structure pointed to by <i>buf</i> :<br><pre>shm_perm.uid shm_perm.gid shm_perm.mode /* only low 9 bits */</pre> <p>This <i>cmd</i> can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of <i>shm_perm.uid</i> in the data structure associated with <i>shmid</i>.</p> |
| IPC_RMID | Remove the shared memory identifier specified by <i>shmid</i> from the system and destroy the shared memory segment and data structure associated with it. This <i>cmd</i> can only be executed by a process that has an effective user ID equal to either that of super-user or to the value of <i>shm_perm.uid</i> in the data structure associated with <i>shmid</i> .                                                                                               |

*Shmctl* will fail if one or more of the following are true:

- |          |                                                                                                                                                                                                                                                           |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [EINVAL] | <i>Shmid</i> is not a valid shared memory identifier.                                                                                                                                                                                                     |
| [EINVAL] | <i>Cmd</i> is not a valid command.                                                                                                                                                                                                                        |
| [EACCES] | <i>Cmd</i> is equal to <b>IPC_STAT</b> and { <b>READ</b> } operation permission is denied to the calling process [see <i>intro</i> (2)].                                                                                                                  |
| [EPERM]  | <i>Cmd</i> is equal to <b>IPC_RMID</b> or <b>IPC_SET</b> and the effective user ID of the calling process is not equal to that of super-user and it is not equal to the value of <i>shm_perm.uid</i> in the data structure associated with <i>shmid</i> . |
| [EFAULT] | <i>Buf</i> points to an illegal address.                                                                                                                                                                                                                  |

## RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

## SEE ALSO

*intro*(2), *shmget*(2), *shmop*(2).

## SHMGET(2)

### NAME

shmget — get shared memory segment

### SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>

int shmget (key, size, shmflg)
key_t key;
int size, shmflg;
```

### DESCRIPTION

*Shmget* returns the shared memory identifier associated with *key*.

A shared memory identifier and associated data structure and shared memory segment of size *size* bytes [see *intro(2)*] are created for *key* if one of the following are true:

*Key* is equal to `IPC_PRIVATE`.

*Key* does not already have a shared memory identifier associated with it, and (*shmflg* & `IPC_CREAT`) is "true".

Upon creation, the data structure associated with the new shared memory identifier is initialized as follows:

`Shm_perm.cuid`, `shm_perm.uid`, `shm_perm.cgid`, and `shm_perm.gid` are set equal to the effective user ID and effective group ID, respectively, of the calling process.

The low-order 9 bits of `shm_perm.mode` are set equal to the low-order 9 bits of *shmflg*. `Shm_segsz` is set equal to the value of *size*.

`Shm_lpid`, `shm_nattch`, `shm_atime`, and `shm_dtime` are set equal to 0.

`Shm_ctime` is set equal to the current time.

*Shmget* will fail if one or more of the following are true:

- |          |                                                                                                                                                                              |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| [EINVAL] | <i>Size</i> is less than the system-imposed minimum or greater than the system-imposed maximum.                                                                              |
| [EACCES] | A shared memory identifier exists for <i>key</i> but operation permission [see <i>intro(2)</i> ] as specified by the low-order 9 bits of <i>shmflg</i> would not be granted. |
| [EINVAL] | A shared memory identifier exists for <i>key</i> but the size of the segment associated with it is less than <i>size</i> and <i>size</i> is not equal to zero.               |
| [ENOENT] | A shared memory identifier does not exist for <i>key</i> and ( <i>shmflg</i> & <code>IPC_CREAT</code> ) is "false".                                                          |
| [ENOSPC] | A shared memory identifier is to be created but the system-imposed limit on the maximum number of allowed shared memory identifiers system wide would be exceeded.           |
| [ENOMEM] | A shared memory identifier and associated shared memory segment are to be created but the amount of available physical memory is not sufficient to fill the request.         |
| [EEXIST] | A shared memory identifier exists for <i>key</i> but ( ( <i>shmflg</i> & <code>IPC_CREAT</code> ) and ( <i>shmflg</i> & <code>IPC_EXCL</code> ) ) is "true".                 |

**RETURN VALUE**

Upon successful completion, a non-negative integer, namely a shared memory identifier is returned. Otherwise, a value of  $-1$  is returned and *errno* is set to indicate the error.

**SEE ALSO**

intro(2), shmctl(2), shmop(2).

S-2

## SHMOP(2)

### NAME

shmop – shared memory operations

### SYNOPSIS

```
#include <sys/types.h>
#include <sys/ipc.h>
#include <sys/shm.h>

char *shmat (shmid, shmaddr, shmflg)
int shmid;
char *shmaddr
int shmflg;

int shmdt (shmaddr)
char *shmaddr
```

### DESCRIPTION

*Shmat* attaches the shared memory segment associated with the shared memory identifier specified by *shmid* to the data segment of the calling process. The segment is attached at the address specified by one of the following criteria:

If *shmaddr* is equal to zero, the segment is attached at the first available address as selected by the system.

If *shmaddr* is not equal to zero and (*shmflg* & SHM\_RND) is “true”, the segment is attached at the address given by (*shmaddr* - (*shmaddr* modulus SHMLBA)).

If *shmaddr* is not equal to zero and (*shmflg* & SHM\_RND) is “false”, the segment is attached at the address given by *shmaddr*.

The segment is attached for reading if (*shmflg* & SHM\_RDONLY) is “true” (READ), otherwise it is attached for reading and writing (READ/WRITE).

*Shmat* will fail and not attach the shared memory segment if one or more of the following are true:

- [EINVAL] *Shmid* is not a valid shared memory identifier.
- [EACCES] Operation permission is denied to the calling process [see *intro(2)*].
- [ENOMEM] The available data space is not large enough to accommodate the shared memory segment.
- [EINVAL] *Shmaddr* is not equal to zero, and the value of (*shmaddr* - (*shmaddr* modulus SHMLBA)) is an illegal address.
- [EINVAL] *Shmaddr* is not equal to zero, (*shmflg* & SHM\_RND) is “false”, and the value of *shmaddr* is an illegal address.
- [EMFILE] The number of shared memory segments attached to the calling process would exceed the system-imposed limit.
- [EINVAL] *Shmdt* detaches from the calling process’s data segment the shared memory segment located at the address specified by *shmaddr*.
- [EINVAL] *Shmdt* will fail and not detach the shared memory segment if *shmaddr* is not the data segment start address of a shared memory segment.



**RETURN VALUES**

Upon successful completion, the return value is as follows:

*Shmat* returns the data segment start address of the attached shared memory segment.

*Shmdt* returns a value of 0.

Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

`exec(2)`, `exit(2)`, `fork(2)`, `intro(2)`, `shmctl(2)`, `shmget(2)`.

## SIGNAL(2)

### NAME

signal — specify what to do upon receipt of a signal

### SYNOPSIS

```
#include <signal.h>
int (*signal (sig, func))()
int sig;
void (*func)();
```

### DESCRIPTION

*Signal* allows the calling process to choose one of three ways in which it is possible to handle the receipt of a specific signal. *Sig* specifies the signal and *func* specifies the choice.

*Sig* can be assigned any one of the following except SIGKILL:

SIGHUP	01	hangup
SIGINT	02	interrupt
SIGQUIT	03*	quit
SIGILL	04*	illegal instruction (not reset when caught)
SIGTRAP	05*	trace trap (not reset when caught)
SIGIOT	06*	IOT instruction
SIGEMT	07*	EMT instruction
SIGFPE	08*	floating point exception
SIGKILL	09	kill (cannot be caught or ignored)
SIGBUS	10*	bus error
SIGSEGV	11*	segmentation violation
SIGSYS	12*	bad argument to system call
SIGPIPE	13	write on a pipe with no one to read it
SIGALRM	14	alarm clock
SIGTERM	15	software termination signal
SIGUSR1	16	user-defined signal 1
SIGUSR2	17	user-defined signal 2
SIGCLD	18	death of a child (see <i>WARNING</i> below)
SIGPWR	19	power fail (see <i>WARNING</i> below)

See below for the significance of the asterisk (\*) in the above list.

*Func* is assigned one of three values: SIG\_DFL, SIG\_IGN, or a *function address*. The actions prescribed by these values are as follows:

**SIG\_DFL** — terminate process upon receipt of a signal

Upon receipt of the signal *sig*, the receiving process is to be terminated with all of the consequences outlined in *exit(2)*. In addition a “core image” will be made in the current working directory of the receiving process if *sig* is one for which an asterisk appears in the above list *and* the following conditions are met:

The effective user ID and the real user ID of the receiving process are equal.

An ordinary file named **core** exists and is writable or can be created. If the file must be created, it will have the following properties:

a mode of 0666 modified by the file creation mask [see *umask(2)*]

a file owner ID that is the same as the effective user ID of the receiving process.

a file group ID that is the same as the effective group ID of the receiving process

**SIG\_IGN** - ignore signal

The signal *sig* is to be ignored.

Note: the signal SIGKILL cannot be ignored.

**function address** - catch signal

Upon receipt of the signal *sig*, the receiving process is to execute the signal-catching function pointed to by *func*. The signal number *sig* will be passed as the only argument to the signal-catching function. Additional arguments are passed to the signal-catching function for hardware-generated signals. Before entering the signal-catching function, the value of *func* for the caught signal will be set to SIG\_DFL unless the signal is SIGILL, SIGTRAP, or SIGPWR.

Upon return from the signal-catching function, the receiving process will resume execution at the point it was interrupted.

When a signal that is to be caught occurs during a *read*, a *write*, an *open*, or an *ioctl* system call on a slow device (like a terminal; but not a file), during a *pause* system call, or during a *wait* system call that does not return immediately due to the existence of a previously stopped or zombie process, the signal catching function will be executed and then the interrupted system call may return a -1 to the calling process with *errno* set to EINTR.

Note: The signal SIGKILL cannot be caught.

A call to *signal* cancels a pending signal *sig* except for a pending SIGKILL signal.

*Signal* will fail if *sig* is an illegal signal number, including SIGKILL. [EINVAL]

**RETURN VALUE**

Upon successful completion, *signal* returns the previous value of *func* for the specified signal *sig*. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

kill(2), pause(2), prctl(2), wait(2), setjmp(3C),  
kill(1) in the Runtime System manual.

**WARNING**

Two other signals that behave differently than the signals described above exist in this release of the system; they are:

SIGCLD	18	death of a child (reset when caught)
SIGPWR	19	power fail (not reset when caught)

There is no guarantee that, in future releases of the UNIX system, these signals will continue to behave as described below; they are included only for compatibility with other versions of the UNIX system. Their use in new programs is strongly discouraged.

For these signals, *func* is assigned one of three values: SIG\_DFL, SIG\_IGN, or a *function address*. The actions prescribed by these values are as follows:

**SIG\_DFL** - ignore signal

The signal is to be ignored.

**SIG\_IGN** - ignore signal

The signal is to be ignored. Also, if *sig* is SIGCLD, the calling process's child processes will not create zombie processes when they terminate; see *exit*(2).

## SIGNAL(2)

### *function address* - catch signal

If the signal is **SIGPWR**, the action to be taken is the same as that described above for *func* equal to *function address*. The same is true if the signal is **SIGCLD** except, that while the process is executing the signal-catching function, any received **SIGCLD** signals will be queued and the signal-catching function will be continually reentered until the queue is empty.

The **SIGCLD** affects two other system calls [*wait(2)*, and *exit(2)*] in the following ways:

*wait*     If the *func* value of **SIGCLD** is set to **SIG\_IGN** and a *wait* is executed, the *wait* will block until all of the calling process's child processes terminate; it will then return a value of **-1** with *errno* set to **ECHILD**.

*exit*     If in the exiting process's parent process the *func* value of **SIGCLD** is set to **SIG\_IGN**, the exiting process will not create a zombie process.

When processing a pipeline, the shell makes the last process in the pipeline the parent of the preceding processes. A process that may be piped into in this manner (and thus become the parent of other processes) should take care not to set **SIGCLD** to be caught.

## NAME

stat, fstat — get file status

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/stat.h>

int stat (path, buf)
char *path;
struct stat *buf;

int fstat (fildes, buf)
int fildes;
struct stat *buf;
```

## DESCRIPTION

*Path* points to a path name naming a file. Read, write, or execute permission of the named file is not required, but all directories listed in the path name leading to the file must be searchable. *Stat* obtains information about the named file.

Similarly, *fstat* obtains information about an open file known by the file descriptor *fildes*, obtained from a successful *open*, *creat*, *dup*, *fcntl*, or *pipe* system call.

*Buf* is a pointer to a *stat* structure into which information is placed concerning the file.

The contents of the structure pointed to by *buf* include the following members:

```
ushort  st_mode; /* File mode; see mknod(2) */
ino_t   st_ino; /* Inode number */
dev_t   st_dev; /* ID of device containing */
          /* a directory entry for this file */
dev_t   st_rdev; /* ID of device */
          /* This entry is defined only for */
          /* character special or block special files */
short   st_nlink; /* Number of links */
ushort  st_uid; /* User ID of the file's owner */
ushort  st_gid; /* Group ID of the file's group */
off_t   st_size; /* File size in bytes */
time_t  st_atime; /* Time of last access */
time_t  st_mtime; /* Time of last data modification */
time_t  st_ctime; /* Time of last file status change */
          /* Times measured in seconds since */
          /* 00:00:00 GMT, Jan. 1, 1970 */

st_atime Time when file data was last accessed. Changed by the following
system calls: creat(2), mknod(2), pipe(2), utime(2), and read(2).

st_mtime Time when data was last modified. Changed by the following sys-
tem calls: creat(2), mknod(2), pipe(2), utime(2), and write(2).

st_ctime Time when file status was last changed. Changed by the following
system calls: chmod(2), chown(2), creat(2), link(2), mknod(2),
pipe(2), unlink(2), utime(2), and write(2).
```

*Stat* will fail if one or more of the following are true:

```
[ENOTDIR] A component of the path prefix is not a directory.
[ENOENT] The named file does not exist.
[EACCES] Search permission is denied for a component of the path
prefix.
```

## STAT(2)

[EFAULT] *Buf* or *path* points to an invalid address.

*Fstat* will fail if one or more of the following are true:

[EBADF] *Fildes* is not a valid open file descriptor.

[EFAULT] *Buf* points to an invalid address.

### RETURN VALUE

Upon successful completion a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

### SEE ALSO

`chmod(2)`, `chown(2)`, `creat(2)`, `link(2)`, `mknod(2)`, `pipe(2)`, `read(2)`, `time(2)`, `unlink(2)`, `utime(2)`, `write(2)`.

NAME  
stime — set time

SYNOPSIS  
int stime (tp)  
long \*tp;

DESCRIPTION  
*Stime* sets the system's idea of the time and date. *Tp* points to the value of time as measured in seconds from 00:00:00 GMT January 1, 1970.

[EPERM] *Stime* will fail if the effective user ID of the calling process is not super-user.

RETURN VALUE  
Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

SEE ALSO  
time(2).

S-2

## **SYNC(2)**

### **NAME**

`sync` — update super-block

### **SYNOPSIS**

`void sync ( )`

### **DESCRIPTION**

*Sync* causes all information in memory that should be on disk to be written out. This includes modified super blocks, modified i-nodes, and delayed block I/O.

It should be used by programs which examine a file system, for example *fsck*, *df*, etc. It is mandatory before a boot.

The writing, although scheduled, is not necessarily complete upon return from *sync*.



## NAME

*time* — get time

## SYNOPSIS

**long time ((long \*) 0)**

**long time (tloc)**

**long \*tloc;**

## DESCRIPTION

*Time* returns the value of time in seconds since 00:00:00 GMT, January 1, 1970.

If *tloc* (taken as an integer) is non-zero, the return value is also stored in the location to which *tloc* points.

[DEFAULT] *Time* will fail if *tloc* points to an illegal address.

## RETURN VALUE

Upon successful completion, *time* returns the value of time. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

## SEE ALSO

*stime*(2).

## TIMES(2)

### NAME

*times* — get process and child process times

### SYNOPSIS

```
#include <sys/types.h>
#include <sys/times.h>

long times (buffer)
struct tms *buffer;
```

### DESCRIPTION

*Times* fills the structure pointed to by *buffer* with time-accounting information. The following are the contents of this structure:

```
struct tms {
    time_t tms_utime;
    time_t tms_stime;
    time_t tms_cutime;
    time_t tms_cstime;
};
```

This information comes from the calling process and each of its terminated child processes for which it has executed a *wait*. All times are in 60ths of a second on DEC and Intel processors, 100ths of a second on AT&T processors.

*Tms\_utime* is the CPU time used while executing instructions in the user space of the calling process.

*Tms\_stime* is the CPU time used by the system on behalf of the calling process.

*Tms\_cutime* is the sum of the *tms\_utimes* and *tms\_cutimes* of the child processes.

*Tms\_cstime* is the sum of the *tms\_stimes* and *tms\_cstimes* of the child processes.

[EFAULT] *Times* will fail if *buffer* points to an illegal address.

### RETURN VALUE

Upon successful completion, *times* returns the elapsed real time, in 60ths (100ths) of a second, since an arbitrary point in the past (e.g., system start-up time). This point does not change from one invocation of *times* to another. If *times* fails, a *-1* is returned and *errno* is set to indicate the error.

### SEE ALSO

*exec(2)*, *fork(2)*, *time(2)*, *wait(2)*.

**NAME**  
uadmin - administrative control

**SYNOPSIS**  
**#include <sys/uadmin.h>**  
**int uadmin (cmd, fcn, mdep)**  
**int cmd, fcn, mdep;**

**DESCRIPTION**

*Uadmin* provides control for basic administrative functions. This system call is tightly coupled to the system administrative procedures and is not intended for general use. The argument *mdep* is provided for machine-dependent use and is not defined here.

The commands available as specified by *cmd* are:

**A\_SHUTDOWN** The system is shutdown. All user processes are killed, the buffer cache is flushed, and the root file system is unmounted. The action to be taken after the system is shutdown is specified by *fcn*. The functions are generic; on specific machines the hardware capabilities will vary.

**AD\_HALT** Halt the processor and turn off power.

**AD\_BOOT** Reboot the system, use /unix.

**AD\_IBOOT** Interactive reboot, prompt for system name.

**A\_REBOOT** The system stops immediately without any further processing. The action to be taken next is specified by *fcn* as above.

**A\_REMOUNT** The root file system is mounted again after having been fixed. This should only be used during the startup process.

*Uadmin* will fail if any of the following are true:

[EPERM] The effective user ID is not super-user.

**DIAGNOSTICS**

Upon successful completion, the value returned depends on *cmd* as follows:

<b>A_SHUTDOWN</b>	Never returns.
<b>A_REBOOT</b>	Never returns.
<b>A_REMOUNT</b>	0

Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

This page intentionally left blank.

**NAME**

ulimit — get and set user limits

**SYNOPSIS**

long ulimit (cmd, newlimit)

int cmd;

long newlimit;

**DESCRIPTION**

This function provides for control over process limits. The *cmd* values available are:

- 1 Get the file size limit of the process. The limit is in units of 512-byte blocks and is inherited by child processes. Files of any size can be read.
- 2 Set the file size limit of the process to the value of *newlimit*. Any process may decrease this limit, but only a process with an effective user ID of super-user may increase the limit. *Ulimit* will fail and the limit will be unchanged if a process with an effective user ID other than super-user attempts to increase its file size limit. [EPERM]
- 3 Get the maximum possible break value. See *brk(2)*.

**RETURN VALUE**

Upon successful completion, a non-negative value is returned. Otherwise, a value of  $-1$  is returned and *errno* is set to indicate the error.

**SEE ALSO**

brk(2), write(2).

## UMASK(2)

### NAME

umask — set and get file creation mask

### SYNOPSIS

```
int umask (cmask)
int cmask;
```

### DESCRIPTION

*Umask* sets the process's file mode creation mask to *cmask* and returns the previous value of the mask. Only the low-order 9 bits of *cmask* and the file mode creation mask are used.

### RETURN VALUE

The previous value of the file mode creation mask is returned.

### SEE ALSO

chmod(2), creat(2), mknod(2), open(2),  
mkdir(1), sh(1) in the Runtime System manual.

**NAME**

mount — unmount a file system

**SYNOPSIS**

```
int mount (spec)
char *spec;
```

**DESCRIPTION**

*Umount* requests that a previously mounted file system contained on the block special device identified by *spec* be unmounted. *Spec* is a pointer to a path name. After unmounting the file system, the directory upon which the file system was mounted reverts to its ordinary interpretation.

*Umount* may be invoked only by the super-user.

*Umount* will fail if one or more of the following are true:

[EPERM]	The process's effective user ID is not super-user.
[ENXIO]	<i>Spec</i> does not exist.
[ENOTBLK]	<i>Spec</i> is not a block special device.
[EINVAL]	<i>Spec</i> is not mounted.
[EBUSY]	A file on <i>spec</i> is busy.
[EFAULT]	<i>Spec</i> points to an illegal address.

**RETURN VALUE**

Upon successful completion a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

mount(2).

## UNAME(2)

### NAME

uname — get name of current UNIX system

### SYNOPSIS

```
#include <sys/utsname.h>
int  uname (name)
struct utsname *name;
```

### DESCRIPTION

*Uname* stores information identifying the current UNIX system in the structure pointed to by *name*.

*Uname* uses the structure defined in `<sys/utsname.h>` whose members are:

```
char  sysname[9];
char  nodename[9];
char  release[9];
char  version[9];
char  machine[9];
```

*Uname* returns a null-terminated character string naming the current UNIX system in the character array *sysname*. Similarly, *nodename* contains the name that the system is known by on a communications network. *Release* and *version* further identify the operating system. *Machine* contains a standard name that identifies the hardware that the UNIX system is running on.

[EFAULT] *Uname* will fail if *name* points to an invalid address.

### RETURN VALUE

Upon successful completion, a non-negative value is returned. Otherwise, `-1` is returned and *errno* is set to indicate the error.

### SEE ALSO

uname(1) in the Runtime System manual.



## NAME

unlink — remove directory entry

## SYNOPSIS

```
int unlink (path)
char *path;
```

## DESCRIPTION

*Unlink* removes the directory entry named by the path name pointed to be *path*.

The named file is unlinked unless one or more of the following are true:

- [ENOTDIR]     A component of the path prefix is not a directory.
- [ENOENT]     The named file does not exist.
- [EACCES]     Search permission is denied for a component of the path prefix.
- [EACCES]     Write permission is denied on the directory containing the link to be removed.
- [EPERM]     The named file is a directory and the effective user ID of the process is not super-user.
- [EBUSY]     The entry to be unlinked is the mount point for a mounted file system.
- [ETXTBSY]    The entry to be unlinked is the last link to a pure procedure (shared text) file that is being executed.
- [EROFS]     The directory entry to be unlinked is part of a read-only file system.
- [EFAULT]     *Path* points outside the process's allocated address space.

When all links to a file have been removed and no process has the file open, the space occupied by the file is freed and the file ceases to exist. If one or more processes have the file open when the last link is removed, the removal is postponed until all references to the file have been closed.

## RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

## SEE ALSO

*close(2)*, *link(2)*, *open(2)*,  
*rm(1)* in the Runtime System manual.

## USTAT(2)

### NAME

ustat — get file system statistics

### SYNOPSIS

```
#include <sys/types.h>
#include <ustat.h>
```

```
int ustat (dev, buf)
int dev;
struct ustat *buf;
```

### DESCRIPTION

*Ustat* returns information about a mounted file system. *Dev* is a device number identifying a device containing a mounted file system. *Buf* is a pointer to a *ustat* structure that includes the following elements:

```
daddr_t f_tfree;      /* Total free blocks */
ino_t    f_tinode;    /* Number of free inodes */
char     f_fname[6];  /* Filsys name */
char     f_fpack[6];  /* Filsys pack name */
```

*Ustat* will fail if one or more of the following are true:

[EINVAL] *Dev* is not the device number of a device containing a mounted file system.

[EFAULT] *Buf* points outside the process's allocated address space.

### RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

### SEE ALSO

stat(2), fs(4).

**NAME**

`utime` - set file access and modification times

**SYNOPSIS**

```
#include <sys/types.h>
int utime (path, times)
char *path;
struct utimbuf *times;
```

**DESCRIPTION**

*Path* points to a path name naming a file. *Utime* sets the access and modification times of the named file.

If *times* is NULL, the access and modification times of the file are set to the current time. A process must be the owner of the file or have write permission to use *utime* in this manner.

If *times* is not NULL, *times* is interpreted as a pointer to a *utimbuf* structure and the access and modification times are set to the values contained in the designated structure. Only the owner of the file or the super-user may use *utime* this way.

The times in the following structure are measured in seconds since 00:00:00 GMT, Jan. 1, 1970.

```
struct utimbuf {
    time_t actime; /* access time */
    time_t modtime; /* modification time */
};
```

*Utime* will fail if one or more of the following are true:

- [ENOENT] The named file does not exist.
- [ENOTDIR] A component of the path prefix is not a directory.
- [EACCES] Search permission is denied by a component of the path prefix.
- [EPERM] The effective user ID is not super-user and not the owner of the file and *times* is not NULL.
- [EACCES] The effective user ID is not super-user and not the owner of the file and *times* is NULL and write access is denied.
- [EROFS] The file system containing the file is mounted read-only.
- [EFAULT] *Times* is not NULL and points outside the process's allocated address space.
- [EFAULT] *Path* points outside the process's allocated address space.

**RETURN VALUE**

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

**SEE ALSO**

`stat(2)`.

## WAIT(2)

### NAME

wait — wait for child process to stop or terminate

### SYNOPSIS

```
int wait (stat_loc)
int *stat_loc;
int wait ((int *)0)
```

### DESCRIPTION

*Wait* suspends the calling process until one of the immediate children terminates or until a child that is being traced stops, because it has hit a break point. The *wait* system call will return prematurely if a signal is received and if a child process stopped or terminated prior to the call on *wait*, return is immediate.

If *stat\_loc* (taken as an integer) is non-zero, 16 bits of information called status are stored in the low order 16 bits of the location pointed to by *stat\_loc*. *Status* can be used to differentiate between stopped and terminated child processes and if the child process terminated, status identifies the cause of termination and passes useful information to the parent. This is accomplished in the following manner:

If the child process stopped, the high order 8 bits of status will contain the number of the signal that caused the process to stop and the low order 8 bits will be set equal to 0177.

If the child process terminated due to an *exit* call, the low order 8 bits of status will be zero and the high order 8 bits will contain the low order 8 bits of the argument that the child process passed to *exit*; see *exit(2)*.

If the child process terminated due to a signal, the high order 8 bits of status will be zero and the low order 8 bits will contain the number of the signal that caused the termination. In addition, if the low order seventh bit (i.e., bit 200) is set, a "core image" will have been produced; see *signal(2)*.

If a parent process terminates without waiting for its child processes to terminate, the parent process ID of each child process is set to 1. This means the initialization process inherits the child processes; see *intro(2)*.

*Wait* will fail and return immediately if one or more of the following are true:

- [ECHILD] The calling process has no existing unwaited-for child processes.
- [EFAULT] *Stat\_loc* points to an illegal address.

### RETURN VALUE

If *wait* returns due to the receipt of a signal, a value of  $-1$  is returned to the calling process and *errno* is set to EINTR. If *wait* returns due to a stopped or terminated child process, the process ID of the child is returned to the calling process. Otherwise, a value of  $-1$  is returned and *errno* is set to indicate the error.

### SEE ALSO

*exec(2)*, *exit(2)*, *fork(2)*, *intro(2)*, *pause(2)*, *ptrace(2)*, *signal(2)*.

### WARNING

See *WARNING* in *signal(2)*.

## NAME

write - write on a file

## SYNOPSIS

```
int write (fildes, buf, nbytes)
int fildes;
char *buf;
unsigned nbytes;
```

## DESCRIPTION

*Fildes* is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, or *pipe* system call.

*Write* attempts to write *nbyte* bytes from the buffer pointed to by *buf* to the file associated with the *fildes*.

On devices capable of seeking, the actual writing of data proceeds from the position in the file indicated by the file pointer. Upon return from *write*, the file pointer is incremented by the number of bytes actually written.

On devices incapable of seeking, writing always takes place starting at the current position. The value of a file pointer associated with such a device is undefined.

If the `O_APPEND` flag of the file status flags is set, the file pointer will be set to the end of the file prior to each write.

For regular files, if the `O_SYNC` flag of the file status flags is set, the write will not return until both the file data and file status are physically updated. This function is for special applications that require extra reliability at the cost of performance. Also, for block special files, if this flag is set, the write will not return until the data has been physically updated.

*Write* will fail and the file pointer will remain unchanged if one or more of the following are true:

[EBADF] *Fildes* is not a valid file descriptor open for writing.

[EPIPE and SIGPIPE signal]

An attempt is made to write to a pipe that is not open for reading by any process.

[EFBIG]

An attempt was made to write a file that exceeds the process's file size limit or the maximum file size. See *ulimit(2)*.

[EFAULT]

*Buf* points outside the process's allocated address space.

[EINTR]

A signal was caught during the *write* system call.

If a *write* requests that more bytes be written than there is room for (e.g., the *ulimit* [see *ulimit(2)*] or the physical end of a medium), only as many bytes as there is room for will be written. For example, suppose there is space for 20 bytes more in a file before reaching a limit. A write of 512 bytes will return 20. The next write of a non-zero number of bytes will give a failure return (except as noted below).

If the file being written is a pipe (or FIFO) and the `O_NDELAY` flag of the file flag word is set, then write to a full pipe (or FIFO) will return a count of 0. Otherwise (`O_NDELAY` clear), writes to a full pipe (or FIFO) will block until space becomes available.

## RETURN VALUE

Upon successful completion, the number of bytes actually written is returned. Otherwise, -1 is returned and *errno* is set to indicate the error.

## **WRITE(2)**

### **SEE ALSO**

`creat(2)`, `dup(2)`, `fcntl(2)`, `lseek(2)`, `open(2)`, `pipe(2)`, `ulimit(2)`.

## TABLE OF CONTENTS OF SUBROUTINES

### 3. Subroutines and Libraries

intro.....	introduction to subroutines and libraries
a64l.....	convert between long integer and base-64 ASCII string
abort.....	generate IOT fault
abs.....	return integer absolute value
bsearch.....	binary search a sorted table
clock.....	report CUP time used
conv.....	translate characters
crypt.....	a one-way hashing encryption algorithm
ctermid.....	generate file name for terminal
ctime.....	convert date and time to string
ctype.....	classify characters
cusetid.....	get character login name of the user
dial.....	establish an out-going terminal line connection
drand48.....	generate uniformly distributed pseudo-random numbers
ecvt.....	convert floating-point number to string
end.....	last locations in program
fclose.....	close or flush a stream
ferror.....	stream status inquiries
fopen.....	open a stream
fread.....	binary input/output
frexp.....	manipulate parts of floating-point numbers
fseek.....	reposition a file pointer in a stream
ftw.....	walk a file tree
getc.....	get character or word from a stream
getcwd.....	get pathname of current working directory
getenv.....	return value for environment name
getgrent.....	get group file entry
getlogin.....	get login name
getopt.....	get option letter from argument vector
getpass.....	read a password
getpw.....	get name from UID
getpwent.....	get password file entry
gets.....	get a string from a stream
getut.....	access utmp file entry
hsearch.....	manage hash search tables
lockf.....	record locking on file
l3tol.....	convert between 3-byte integers and long integers
lsearch.....	linear search and update
malloc.....	main memory allocator
memory.....	memory operations
mktemp.....	make a unique file name
monitor.....	prepare execution profile
nlist.....	get entries from name list
perror.....	system error messages
popen.....	initiate pipe to/from a process
printf.....	print formatted output
putc.....	put character or word on a stream
putenv.....	change or add value to environment
putpwent.....	write password file entry
puts.....	put a string on a stream
qsort.....	quicker sort
rand.....	simple random-number generator
scanf.....	convert formatted input
setbuf.....	assign buffering to a stream
setjmp.....	non-local goto
sleep.....	suspend execution for interval
signal.....	software signals

stdio.....	standard buffered input/output package
stdipc.....	standard interprocess communication package
string.....	string operations
strtod.....	convert string to double-precision number
strtol.....	convert string to integer
swab.....	swap bytes
system.....	issue a shell command
tmpfile.....	create a temporary file
tmpnam.....	create a name for a temporary file
tsearch.....	manage binary search trees
tyname.....	find name of a terminal
ttyslot.....	find the slot in the utmp file of the current user
ungetc.....	push character back into input stream
vprintf.....	print formatted output of args argument list

## TABLE OF CONTENTS OF MATH SUBROUTINES

### 3M. Math Subroutines and Libraries

bessel.....	Bessel functions
erf.....	error function and complementary error function
exp.....	exponential, logarithm, power, square root functions
floor.....	floor, ceiling, remainder absolute value functions
gamma.....	log gamma function
hypot.....	Euclidean distance function
matherr.....	error-handling function
sinh.....	hyperbolic functions
trig.....	trigonometric functions



## NAME

intro — introduction to subroutines and libraries

## SYNOPSIS

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

```
#include <math.h>
```

## DESCRIPTION

This section describes functions found in various libraries, other than those functions that directly invoke UNIX system primitives, which are described in Section 2 of this volume. Certain major collections are identified by a letter after the section number:

- (3C) These functions, together with those of Section 2 and those marked (3S), constitute the Standard C Library *libc*, which is automatically loaded by the C compiler, *cc*(1). The link editor *ld*(1) searches this library under the *-lc* option. Declarations for some of these functions may be obtained from **#include** files indicated on the appropriate pages.
- (3S) These functions constitute the "standard I/O package" [see *stdio*(3S)]. These functions are in the library *libc*, already mentioned. Declarations for these functions may be obtained from the **#include** file *<stdio.h>*.
- (3M) These functions constitute the Math Library, *libm*. They are automatically loaded as needed by the FORTRAN compiler *f77*(1). They are not automatically loaded by the C compiler, *cc*(1); however, the link editor searches this library under the *-lm* option. Declarations for these functions may be obtained from the **#include** file *<math.h>*. Several generally useful mathematical constants are also defined there [see *math*(5)].
- (3X) Various specialized libraries. The files in which these libraries are found are given on the appropriate pages.
- (3F) These functions constitute the FORTRAN intrinsic function library, *libF77*. These functions are automatically available to the FORTRAN programmer and require no special invocation of the compiler.

There are separate library files for use with small, large and huge model programs (see "Files," below). Normally, *ld*(1) automatically selects the correct library file for the memory model you are using. However, if you specify the library file yourself, be sure that it matches the memory model of your program.

## DEFINITIONS

A *character* is any bit pattern able to fit into a byte on the machine. The *null character* is a character with value 0, represented in the C language as `'\0'`. A *character array* is a sequence of characters. A *null-terminated character array* is a sequence of characters, the last of which is the *null character*. A *string* is a designation for a *null-terminated character array*. The *null string* is a character array containing only the null character. A *NULL pointer* is the value that is obtained by casting 0 into a pointer. The C language guarantees that this value will not match that of any legitimate pointer, so many functions that return pointers return it to indicate an error. *NULL* is defined as 0 in *<stdio.h>*; the user can include an appropriate definition if not using *<stdio.h>*.

Many groups of FORTRAN intrinsic functions have *generic* function names that do not require explicit or implicit type declaration. The type of the function will be determined by the type of its argument(s). For example, the generic function *max* will return an integer value if given integer arguments (*max0*), a real value if given real arguments (*amax1*), or a double-precision value if given double-precision arguments (*dmax1*).

## INTRO(3)

### FILES

/lib/small/libc.a  
/lib/large/libc.a  
/lib/huge/libc.a  
/lib/small/libm.a  
/lib/large/libm.a  
/lib/huge/libm.a  
/usr/small/lib/libF77.a  
/usr/large/lib/libF77.a  
/usr/huge/lib/libF77.a

### SEE ALSO

intro(2), stdio(3S), math(5).  
ar(1), cc(1), f77(1), ld(1), lint(1), nm(1) in the Runtime System manual.

### DIAGNOSTICS

Functions in the C and Math Libraries (3C and 3M) may return the conventional values 0 or  $\pm$ HUGE (the largest-magnitude single-precision floating-point numbers; HUGE is defined in the `<math.h>` header file) when the function is undefined for the given arguments or when the value is not representable. In these cases, the external variable `errno` [see *intro*(2)] is set to the value EDOM or ERANGE. As many of the FORTRAN intrinsic functions use the routines found in the Math Library, the same conventions apply.

### WARNING

Many of the functions in the libraries call and/or refer to other functions and external variables described in this section and in section 2 (*System Calls*). If a program inadvertently defines a function or external variable with the same name, the presumed library version of the function or external variable may not be loaded. The `lint(1)` program checker reports name conflicts of this kind as "multiple declarations" of the names in question. Definitions for sections 2, 3C, and 3S are checked automatically. Other definitions can be included by using the `-I` option (for example, `-Im` includes definitions for the Math Library, section 3M). Use of `lint` is highly recommended.

## NAME

*a64l*, *l64a* — convert between long integer and base-64 ASCII string

## SYNOPSIS

```
long a64l (s)
char *s;
char *l64a (l)
long l;
```

## DESCRIPTION

These functions are used to maintain numbers stored in *base-64* ASCII characters. This is a notation by which long integers can be represented by up to six characters; each character represents a "digit" in a radix-64 notation.

The characters used to represent "digits" are . for 0, / for 1, 0 through 9 for 2–11, A through Z for 12–37, and a through z for 38–63.

*A64l* takes a pointer to a null-terminated base-64 representation and returns a corresponding long value. If the string pointed to by *s* contains more than six characters, *a64l* will use the first six.

*L64a* takes a long argument and returns a pointer to the corresponding base-64 representation. If the argument is 0, *l64a* returns a pointer to a null string.

## BUGS

The value returned by *l64a* is a pointer into a static buffer, the contents of which are overwritten by each call.

5-3

# ABORT(3C)

## NAME

`abort` — generate an IOT fault

## SYNOPSIS

```
int abort ( )
```

## DESCRIPTION

*Abort* first closes all open files if possible, then causes an IOT signal to be sent to the process. This usually results in termination with a core dump.

It is possible for *abort* to return control if SIGIOT is caught or ignored, in which case the value returned is that of the *kill(2)* system call.

## SEE ALSO

*exit(2)*, *kill(2)*, *signal(2)*.

*sdb(1)* in the Runtime System manual.

*val*.

## DIAGNOSTICS

If SIGIOT is neither caught nor ignored, and the current directory is writable, a core dump is produced and the message “*abort — core dumped*” is written by the shell.

**NAME**  
abs — return integer absolute value

**SYNOPSIS**  
int abs (i)  
int i;

**DESCRIPTION**  
*Abs* returns the absolute value of its integer operand.

**BUGS**  
In two's-complement representation, the absolute value of the negative integer with largest magnitude is undefined. Some implementations trap this error, but others simply ignore it.

**SEE ALSO**  
floor(3M).

## BSEARCH(3C)

### NAME

bsearch — binary search a sorted table

### SYNOPSIS

```
#include <search.h>
```

```
char *bsearch ((char *) key, (char *) base, nel, sizeof (*key), compar)
unsigned nel;
int (*compar)( );
```

### DESCRIPTION

*Bsearch* is a binary search routine generalized from Knuth (6.2.1) Algorithm B. It returns a pointer into a table indicating where a datum may be found. The table must be previously sorted in increasing order according to a provided comparison function. *Key* points to a datum instance to be sought in the table. *Base* points to the element at the base of the table. *Nel* is the number of elements in the table. *Compar* is the name of the comparison function, which is called with two arguments that point to the elements being compared. The function must return an integer less than, equal to, or greater than zero as accordingly the first argument is to be considered less than, equal to, or greater than the second.

### EXAMPLE

The example below searches a table containing pointers to nodes consisting of a string and its length. The table is ordered alphabetically on the string in the node pointed to by each entry.

This code fragment reads in strings and either finds the corresponding node and prints out the string and its length, or prints an error message.

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

#define TABSIZE      1000

struct node {
    char *string;
    int length;
};
struct node table[TABSIZE]; /* table to be searched */
.
.
.
(
    struct node *node_ptr, node;
    int node_compare( ); /* routine to compare 2 nodes */
    char str_space[20]; /* space to read string into */
    .
    .
    .
    node.string = str_space;
    while (scanf("%s", node.string) != EOF) {
        node_ptr = (struct node *)bsearch((char *)&node,
            (char *)table, TABSIZE,
            sizeof(struct node), node_compare);
        if (node_ptr != NULL) {
            (void)printf("string = %20s, length = %d\n",
                node_ptr->string, node_ptr->length);
        } else {
            (void)printf("not found: %s\n", node.string);
        }
    }
)
```

```

    }
}
/*
   This routine compares two nodes based on an
   alphabetical ordering of the string field.
*/
int
node_compare(node1, node2)
struct node *node1, *node2;
{
    return strcmp(node1->string, node2->string);
}

```

**NOTES**

The pointers to the key and the element at the base of the table should be of type pointer-to-element, and cast to type pointer-to-character.

The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

Although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

**SEE ALSO**

bsearch(3C), lsearch(3C), qsort(3C), tsearch(3C).

**DIAGNOSTICS**

A NULL pointer is returned if the key cannot be found in the table.

## CLOCK(3C)

### NAME

clock — report CPU time used

### SYNOPSIS

**long clock ( )**

### DESCRIPTION

*Clock* returns the amount of CPU time (in microseconds) used since the first call to *clock*. The time reported is the sum of the user and system times of the calling process and its terminated child processes for which it has executed *wait(2)* or *system(3S)*.

The resolution of the clock is 10 milliseconds on AT&T 3B computers, 16.667 milliseconds on Digital Equipment Corporation and Intel processors.

### SEE ALSO

*times(2)*, *wait(2)*, *system(3S)*.

### BUGS

The value returned by *clock* is defined in microseconds for compatibility with systems that have CPU clocks with much higher resolution. Because of this, the value returned will wrap around after accumulating only 2147 seconds of CPU time (about 36 minutes).



## NAME

*toupper*, *tolower*, *\_toupper*, *\_tolower*, *toascii* — translate characters

## SYNOPSIS

```
#include <ctype.h>
int toupper (c)
int c;
int tolower (c)
int c;
int _toupper (c)
int c;
int _tolower (c)
int c;
int toascii (c)
int c;
```

## DESCRIPTION

*Toupper* and *tolower* have as domain the range of *getc*(3S): the integers from -1 through 255. If the argument of *toupper* represents a lowercase letter, the result is the corresponding uppercase letter. If the argument of *tolower* represents an uppercase letter, the result is the corresponding lowercase letter. All other arguments in the domain are returned unchanged.

The macros *\_toupper* and *\_tolower*, are macros that accomplish the same thing as *toupper* and *tolower* but have restricted domains and are faster. *\_toupper* requires a lowercase letter as its argument; its result is the corresponding uppercase letter. The macro *\_tolower* requires an uppercase letter as its argument; its result is the corresponding lowercase letter. Arguments outside the domain cause undefined results.

*Toascii* yields its argument with all bits turned off that are not part of a standard ASCII character; it is intended for compatibility with other systems.

## SEE ALSO

*ctype*(3C), *getc*(3S).

## CRYPT(3C)

### NAME

crypt, encrypt — a one-way hashing encryption algorithm

### SYNOPSIS

```
char *crypt (key, salt)
char *key, *salt;

void encrypt (block)
char *block;
```

### DESCRIPTION

*Crypt* is the password encryption function. It is based on a one-way hashing encryption algorithm with variations intended (among other things) to frustrate use of hardware implementations of a key search.

*Key* is a user's typed password. *Salt* is a two-character string chosen from the set [a-zA-Z0-9./]; this string is used to perturb the hashing algorithm in one of 4096 different ways, after which the password is used as the key to encrypt repeatedly a constant string. The returned value points to the encrypted password. The first two characters are the salt itself.

There is a character array of length 64 containing only the characters with numerical value 0 and 1. When this string is divided into groups of 8, the low-order bit in each group is ignored; this gives a 56-bit key which is set into the machine by *crypt*.

The *encrypt* entry provides (rather primitive) access to the actual hashing algorithm. The argument to the *encrypt* entry is a character array of length 64 containing only the characters with numerical value of 0 and 1. The argument array is modified in place to a similar array representing the bits of the argument after having been subjected to the hashing algorithm using the key set by *crypt*.

### SEE ALSO

getpass(3C), passwd(4),  
login(1), passwd(1) in the Runtime System manual.

### BUGS

The return value points to static data that are overwritten by each call.

**NAME**

*ctermid* — generate file name for terminal

**SYNOPSIS**

```
#include <stdio.h>
char *ctermid (s)
char *s;
```

**DESCRIPTION**

*Ctermid* generates the path name of the controlling terminal for the current process, and stores it in a string.

If *s* is a NULL pointer, the string is stored in an internal static area, the contents of which are overwritten at the next call to *ctermid*, and the address of which is returned. Otherwise, *s* is assumed to point to a character array of at least **L\_ctermid** elements; the path name is placed in this array and the value of *s* is returned. The constant **L\_ctermid** is defined in the *<stdio.h>* header file.

**NOTES**

The difference between *ctermid* and *ttyname*(3C) is that *ttyname* must be handed a file descriptor and returns the actual name of the terminal associated with that file descriptor, while *ctermid* returns a string (*/dev/tty*) that will refer to the terminal if used as a file name. Thus *ttyname* is useful only if the process already has at least one file open to a terminal.

**SEE ALSO**

*ttyname*(3C).

## CTIME(3C)

### NAME

*ctime*, *localtime*, *gmtime*, *asctime*, *tzset* — convert date and time to string

### SYNOPSIS

```
#include <time.h>
char *ctime (clock)
long *clock;

struct tm *localtime (clock)
long *clock;

struct tm *gmtime (clock)
long *clock;

char *asctime (tm)
struct tm *tm;

extern long timezone;
extern int daylight;
extern char *tzname[2];
void tzset ( )
```

### DESCRIPTION

*Ctime* converts a long integer, pointed to by *clock*, representing the time in seconds since 00:00:00 GMT, January 1, 1970, and returns a pointer to a 26-character string in the following form. All the fields have constant width.

```
Sun Sep 16 01:03:52 1973\n\n0
```

*Localtime* and *gmtime* return pointers to “tm” structures, described below. *Localtime* corrects for the time zone and possible Daylight Saving Time; *gmtime* converts directly to Greenwich Mean Time (GMT), which is the time the UNIX system uses.

*Asctime* converts a “tm” structure to a 26-character string, as shown in the above example, and returns a pointer to the string.

Declarations of all the functions and externals, and the “tm” structure, are in the *<time.h>* header file. The structure declaration is:

```
struct tm {
    int tm_sec;        /* seconds (0 - 59) */
    int tm_min;       /* minutes (0 - 59) */
    int tm_hour;      /* hours (0 - 23) */
    int tm_mday;      /* day of month (1 - 31) */
    int tm_mon;       /* month of year (0 - 11) */
    int tm_year;      /* year - 1900 */
    int tm_wday;      /* day of week (Sunday = 0) */
    int tm_yday;      /* day of year (0 - 365) */
    int tm_isdst;
};
```

*Tm\_isdst* is non-zero if Daylight Saving Time is in effect.

The external *long* variable *timezone* contains the difference, in seconds, between GMT and local standard time (in EST, *timezone* is 5\*60\*60); the external variable *daylight* is non-zero if and only if the standard U.S.A. Daylight Saving Time conversion should be applied. The program knows about the peculiarities of this conversion in 1974 and 1975; if necessary, a table for these years can be extended.

If an environment variable named TZ is present, *asctime* uses the contents of the variable to override the default time zone. The value of TZ must be a three-letter time zone name, followed by a number representing the difference between local time and Greenwich Mean Time in hours, followed by an optional three-letter name for a daylight time zone. For example, the setting for New Jersey would be EST5EDT. The effects of setting TZ are thus to change the values of the external variables *timezone* and *daylight*; in addition, the time zone names contained in the external variable

```
char *tzname[2] = { "EST", "EDT" };
```

are set from the environment variable TZ. The function *tzset* sets these external variables from TZ; *tzset* is called by *asctime* and may also be called explicitly by the user.

Note that in most installations, TZ is set by default when the user logs on, to a value in the local */etc/profile* file [see *profile* (4)].

#### SEE ALSO

time(2), getenv(3C), profile(4), environ(5).

#### BUGS

The return values point to static data whose content is overwritten by each call.

## CTYPE(3C)

### NAME

*isalpha*, *isupper*, *islower*, *isdigit*, *isxdigit*, *isalnum*, *isspace*, *ispunct*, *isprint*, *isgraph*, *isctrl*, *isascii* – classify characters

### SYNOPSIS

```
#include <ctype.h>
```

```
int isalpha (c)
```

```
int c;
```

```
...
```

### DESCRIPTION

These macros classify character-coded integer values by table lookup. Each is a predicate returning nonzero for true, zero for false. *isascii* is defined on all integer values; the rest are defined only where *isascii* is true and on the single non-ASCII value EOF [-1 – see *stdio(3S)*].

*isalpha*            *c* is a letter.

*isupper*           *c* is an uppercase letter.

*islower*           *c* is a lowercase letter.

*isdigit*            *c* is a digit [0-9].

*isxdigit*          *c* is a hexadecimal digit [0-9], [A-F] or [a-f].

*isalnum*            *c* is an alphanumeric (letter or digit).

*isspace*            *c* is a space, tab, carriage return, new-line, vertical tab, or form-feed.

*ispunct*            *c* is a punctuation character (neither control nor alphanumeric).

*isprint*            *c* is a printing character, code 040 (space) through 0176 (tilde).

*isgraph*            *c* is a printing character, like *isprint* except false for space.

*isctrl*             *c* is a delete character (0177) or an ordinary control character (less than 040).

*isascii*            *c* is an ASCII character, code less than 0200.

### DIAGNOSTICS

If the argument to any of these macros is not in the domain of the function, the result is undefined.

### SEE ALSO

*stdio(3S)*, *ascii(5)*.

**NAME**

`cuserid` — get character login name of the user

**SYNOPSIS**

```
#include <stdio.h>
char *cuserid (s)
char *s;
```

**DESCRIPTION**

*Cuserid* generates a character-string representation of the login name that the owner of the current process is logged in under. If *s* is a NULL pointer, this representation is generated in an internal static area, the address of which is returned. Otherwise, *s* is assumed to point to an array of at least `L_cuserid` characters; the representation is left in this array. The constant `L_cuserid` is defined in the `<stdio.h>` header file.

**DIAGNOSTICS**

If the login name cannot be found, *cuserid* returns a NULL pointer; if *s* is not a NULL pointer, a null character (`\0`) will be placed at *s[0]*.

**SEE ALSO**

`getlogin(3C)`, `getpwent(3C)`.

**This page intentionally left blank.**



## NAME

dial — dial out on a modem

## SYNOPSIS

```
#include <termio.h>
#include <dial.h>

int dial (call)
CALL call;

void undial (fd)
int fd;

extern int debug ; /* set to extern int nolock ; /* don't
extern int (*lockfn) ( ), (*unlockfn) ( ) ; /*
```

## DESCRIPTION

*Dial* is a modem-independent implementation of the standard System V *Dial* procedure.

*Dial* calls out on a modem or a direct terminal line, and returns a file-descriptor open for read/write.

When finished with the line, the caller must invoke *undial* to gracefully disconnect.

*Dial* reads the uucp *L-devices* file to determine eligible devices, and the *dialinfo* file to determine the modem dial procedure. *Dial* is fully compatible with *cu* and *uucp*.

The definition of CALL in the <dial.h> header file is:

```
typedef struct {
    struct termio *attr; /* Final terminal attributes */
    int baud; /* Baud rate to use after dialing */
    int speed; /* Baud rate to use during dialing */
    char *line; /* TTY device name */
    char *telno; /* Phone number(s) or system name */
    int modem; /* Use modem control on direct lines */
    char *device; /* Where to store device name */
    int dev_len; /* Sizeof (call.device) */
} CALL;
```

The CALL parameters are as follows:

- |       |                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| attr  | If specified, the tty device is set to use the given parity, character size, and baud rate after connect. Before a successful return, the remaining tty attributes—except carrier sense—are set. Carrier sense is controlled exclusively through <i>dialinfo</i> commands. If not given, baud rate defaults to the modem baud rate, and in any case is overridden by <i>baud</i> below, or commands in the <i>dialinfo</i> file. |
| baud  | If specified, the tty device is set to this baud rate after connect. This parameter may be overridden by commands in the <i>dialinfo</i> file.                                                                                                                                                                                                                                                                                   |
| speed | If specified, only entries with this speed in field 4 of <i>L-devices</i> are considered.                                                                                                                                                                                                                                                                                                                                        |
| line  | If specified, only entries with this tty name in field 2 of <i>L-devices</i> are considered.                                                                                                                                                                                                                                                                                                                                     |
| telno | If not specified, the call is to a direct line. Otherwise this is the telephone number to be dialed. Several alternate numbers may be given, separated by commas. <i>L-dialcodes</i> prefix substitution is performed.                                                                                                                                                                                                           |

## DIAL(3C)

- modem** Modem control is used on direct lines when this variable is non-zero, and connect option 2 (command C2) is selected by the *dialinfo* procedure.
- device** If this pointer is non-null, the line device pathname (eg */dev/tty12*) is copied here after a successful dial.
- dev\_len** Length of the array pointed to by *device* above.

If the global variable *lock* is a non-zero, *dial* does not test for or secure the normal uucp lock file, and *undial* does not remove it.

### FILES

- /usr/lib/uucp/dialinfo* Dial procedure data base.  
*/usr/lib/uucp/L-devices* UUCP device entries.  
*/usr/lib/uucp/L-dialcodes* Dialcode substitutions.  
*/usr/lib/uucpL.sys* UUCP system definitions.  
*/usr/spool/uucp/LCK..\** UUCP lock file.  
*/dev/tty\** TTY device names.

### AUTHOR

G ne H. Olson, Quest Research, Burnsville, MN. Daniel M. Frank, Microport Systems, wrote the locking modifications.

### SEE ALSO

- cu(1C)* Calls another system.  
*dialprint(1)* Prints a dialer entry.  
*dialer(1)* Modem dial-out program.  
*uucp(1)* UUCP interface.  
*dialinfo(4)* Dial procedure data base.  
*termio(7)* TTY device information

### DIAGNOSTICS

On failure, a negative error code is returned. Possible error codes—as listed in *<dial.h>* include:

- |         |     |                                            |
|---------|-----|--------------------------------------------|
| INTRPT  | -1  | /* Interrupt during dial */                |
| D_HUNG  | -2  | /* Dialer hung */                          |
| NO_ANS  | -3  | /* Busy or no answer */                    |
| ILL_BD  | -4  | /* Illegal/unknown baud rate */            |
| A_PROB  | -5  | /* Dialinfo(4) configuration error */      |
| L_PROB  | -6  | /* TTY device error */                     |
| NO_Ldv  | -7  | /* L-devices file unreadable */            |
| DV_NT_A | -8  | /* Requested device not available */       |
| DV_NT_K | -9  | /* Requested device unknown */             |
| NO_BD_A | -10 | /* Nothing available at requested speed */ |
| NO_BD_K | -11 | /* No device known at requested speed */   |

If the environment variable *DIALINFO* exists, it specifies a pathname to be used instead of */usr/lib/uucp/dialinfo*.

If *dial* discovers a corrupted or improperly configured file or device, diagnostic information is written to *stderr*.

Debugging output is written to *stderr* when the environment variable *DIALDEBUG* exists, and contains one or more of the characters listed below.

- d Show data and decisions related to the *L-devices* file, plus creation and deletion of the *uucplock* file.
- s Show dialer state definitions and transitions.
- m Show matching of dialer transition strings against the incoming data stream.
- l Show all device operations other than character I/O to the communications device.
- c Show all transmitted and received communication.
- a Show all of the above.

The *dialinfo* file may also contain commands which write user specified diagnostics to *stderr*.

#### WARNING

When internal locking is used, the dial lock file is touched every hour by a routine triggered by the *alarm()*. If you use the alarm signal for any purpose, be very careful to restore both the signal and the alarm clock value when you are done. If you need the alarm on an ongoing basis, you must establish your own locking routines, and inform *dial()* of those routines using the external variables *lockfn* and *unlockfn*. Then it becomes your responsibility to touch the lock file.

## DRAND48(3C)

### NAME

*drand48*, *crand48*, *lrand48*, *nrand48*, *mrand48*, *jrand48*, *srand48*, *seed48*, *lcng48* — generate uniformly distributed pseudo-random numbers

### SYNOPSIS

```
double drand48 ( )
double erand48 (xsubi)
unsigned short xsubil[3];
long lrand48 ( )
long nrand48 (xsubi)
unsigned short xsubil[3];
long mrand48 ( )
long jrand48 (xsubi)
unsigned short xsubil[3];
void srand48 (seedval)
long seedval;
unsigned short *seed48 (seed16v)
unsigned short seed16v[3];
void lcng48 (param)
unsigned short param[7];
```

### DESCRIPTION

This family of functions generates pseudo-random numbers using the well-known linear congruential algorithm and 48-bit integer arithmetic.

Functions *drand48* and *erand48* return non-negative double-precision floating-point values uniformly distributed over the interval [0.0, 1.0).

Functions *lrand48* and *nrand48* return non-negative long integers uniformly distributed over the interval [0,  $2^{31}$ ).

Functions *mrand48* and *jrand48* return signed long integers uniformly distributed over the interval [ $-2^{31}$ ,  $2^{31}$ ).

Functions *srand48*, *seed48* and *lcng48* are initialization entry points, one of which should be invoked before either *drand48*, *lrand48* or *mrand48* is called. (Although it is not recommended practice, constant default initializer values will be supplied automatically if *drand48*, *lrand48* or *mrand48* is called without a prior call to an initialization entry point.) Functions *erand48*, *nrand48* and *jrand48* do not require an initialization entry point to be called first.

All the routines work by generating a sequence of 48-bit integer values,  $X_i$ , according to the linear congruential formula

$$X_{n+1} = (aX_n + c)_{\text{mod } m} \quad n \geq 0.$$

The parameter  $m = 2^{48}$ ; hence 48-bit integer arithmetic is performed. Unless *lcng48* has been invoked, the multiplier value  $a$  and the addend value  $c$  are given by

$$a = 5DEECE66D_{16} = 273673163155_8$$
$$c = B_{16} = 13_8.$$

The value returned by any of the functions *drand48*, *erand48*, *lrand48*, *nrand48*, *mrand48* or *jrand48* is computed by first generating the next 48-bit  $X_i$  in the sequence. Then the appropriate number of bits, according to the type of data item to be returned, are copied from the high-order (leftmost) bits of  $X_i$  and transformed into the returned value.

The functions *drand48*, *lrand48* and *mrnd48* store the last 48-bit  $X_i$  generated in an internal buffer; that is why they must be initialized prior to being invoked. The functions *erand48*, *nrnd48* and *jrnd48* require the calling program to provide storage for the successive  $X_i$  values in the array specified as an argument when the functions are invoked. That is why these routines do not have to be initialized; the calling program merely has to place the desired initial value of  $X_i$  into the array and pass it as an argument. By using different arguments, functions *erand48*, *nrnd48* and *jrnd48* allow separate modules of a large program to generate several *independent* streams of pseudo-random numbers, i.e., the sequence of numbers in each stream will *not* depend upon how many times the routines have been called to generate numbers for the other streams.

The initializer function *srand48* sets the high-order 32 bits of  $X_i$  to the 32 bits contained in its argument. The low-order 16 bits of  $X_i$  are set to the arbitrary value  $330E_{16}$ .

The initializer function *seed48* sets the value of  $X_i$  to the 48-bit value specified in the argument array. In addition, the previous value of  $X_i$  is copied into a 48-bit internal buffer, used only by *seed48*, and a pointer to this buffer is the value returned by *seed48*. This returned pointer, which can just be ignored if not needed, is useful if a program is to be restarted from a given point at some future time — use the pointer to get at and store the last  $X_i$  value, and then use this value to reinitialize via *seed48* when the program is restarted.

The initialization function *lcong48* allows the user to specify the initial  $X_i$ , the multiplier value  $a$ , and the addend value  $c$ . Argument array elements *param[0-2]* specify  $X_i$ , *param[3-5]* specify the multiplier  $a$ , and *param[6]* specifies the 16-bit addend  $c$ . After *lcong48* has been called, a subsequent call to either *srand48* or *seed48* will restore the “standard” multiplier and addend values,  $a$  and  $c$ , specified on the previous page.

#### NOTES

The versions of these routines for the VAX-11 and PDP-11 are coded in assembly language for maximum speed. It requires approximately 80  $\mu$ sec on a VAX-11/780 and 130  $\mu$ sec on a PDP-11/70 to generate one pseudo-random number. On other computers, the routines are coded in portable C. The source code for the portable version can even be used on computers which do not have floating-point arithmetic. In such a situation, functions *drand48* and *erand48* do not exist; instead, they are replaced by the two new functions below.

**long irand48 (m)**  
**unsigned short m;**

**long krand48 (xsubi, m)**  
**unsigned short xsubi[3], m;**

Functions *irand48* and *krand48* return non-negative long integers uniformly distributed over the interval  $[0, m-1]$ .

SEE ALSO  
 rand(3C).

## ECVT(3C)

### NAME

*ecvt*, *fcvt*, *gcvt* — convert floating-point number to string

### SYNOPSIS

```
char *ecvt (value, ndigit, decpt, sign)
double value;
int ndigit, *decpt, *sign;

char *fcvt (value, ndigit, decpt, sign)
double value;
int ndigit, *decpt, *sign;

char *gcvt (value, ndigit, buf)
double value;
int ndigit;
char *buf;
```

### DESCRIPTION

*Ecvt* converts *value* to a null-terminated string of *ndigit* digits and returns a pointer thereto. The high-order digit is non-zero, unless the value is zero. The low-order digit is rounded. The position of the decimal point relative to the beginning of the string is stored indirectly through *decpt* (negative means to the left of the returned digits). The decimal point is not included in the returned string. If the sign of the result is negative, the word pointed to by *sign* is non-zero, otherwise it is zero.

*Fcvt* is identical to *ecvt*, except that the correct digit has been rounded for printf “%f” (FORTRAN F-format) output of the number of digits specified by *ndigit*.

*Gcvt* converts the *value* to a null-terminated string in the array pointed to by *buf* and returns *buf*. It attempts to produce *ndigit* significant digits in FORTRAN F-format if possible, otherwise E-format, ready for printing. A minus sign, if there is one, or a decimal point will be included as part of the returned string. Trailing zeros are suppressed.

### SEE ALSO

printf(3S).

### BUGS

The values returned by *ecvt* and *fcvt* point to a single static data array whose content is overwritten by each call.

## NAME

end, etext, edata — last locations in program

## SYNOPSIS

```
extern end;  
extern void etext();  
extern edata;
```

## DESCRIPTION

These names refer neither to routines nor to locations with interesting contents. The address of *etext* is the first address above the program text, *edata* above the initialized data region, and *end* above the uninitialized data region. The addresses are logical and have the form *selector:offset* in large and huge model.

When execution begins, the program break (the first location beyond the data) coincides with *end*, but the program break may be reset by the routines of *brk(2)*, *malloc(3C)*, standard input/output [*stdio(3S)*], the profile (*-p*) option of *cc(1)*, and so on. In small model programs *sbrk(0)* returns the current value of the program break; in large and huge model programs, it returns the starting address of the next data segment [see *brk(2)*].

## SEE ALSO

*brk(2)*, *malloc(3C)*, *stdio(3S)*,  
*cc(1)* in the Runtime System manual.

## FCLOSE(3S)

### NAME

`fclose`, `fflush` — close or flush a stream

### SYNOPSIS

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

```
int fclose (stream)
```

```
FILE *stream;
```

```
int fflush (stream)
```

```
FILE *stream;
```

### DESCRIPTION

*Fclose* causes any buffered data for the named *stream* to be written out, and the *stream* to be closed.

*Fclose* is performed automatically for all open files upon calling *exit*(2).

*Fflush* causes any buffered data for the named *stream* to be written to that file. The *stream* remains open.

### DIAGNOSTICS

These functions return 0 for success, and EOF if any error (such as trying to write to a file that has not been opened for writing) was detected.

### SEE ALSO

`close`(2), `exit`(2), `fopen`(3S), `setbuf`(3S).



**NAME**

*ferror*, *feof*, *clearerr*, *fileno* — stream status inquiries

**SYNOPSIS**

```
#include <stdio.h>

int ferror (stream)
FILE *stream;

int feof (stream)
FILE *stream;

void clearerr (stream)
FILE *stream;

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

**DESCRIPTION**

*Ferror* returns non-zero when an I/O error has previously occurred reading from or writing to the named *stream*, otherwise zero.

*Feof* returns non-zero when EOF has previously been detected reading the named input *stream*, otherwise zero.

*Clearerr* resets the error indicator and EOF indicator to zero on the named *stream*.

*Fileno* returns the integer file descriptor associated with the named *stream*; see *open(2)*.

**NOTE**

All these functions are implemented as macros; they cannot be declared or redeclared.

**SEE ALSO**

*open(2)*, *fopen(3S)*.

5  
3

# FOPEN(3S)

## NAME

fopen, freopen, fdopen — open a stream

## SYNOPSIS

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

```
FILE *fopen (file-name, type)
```

```
char *file-name, *type;
```

```
FILE *freopen (file-name, type, stream)
```

```
char *file-name, *type;
```

```
FILE *stream;
```

```
FILE *fdopen (fdes, type)
```

```
int fdes;
```

```
char *type;
```

## DESCRIPTION

*Fopen* opens the file named by *file-name* and associates a *stream* with it. *Fopen* returns a pointer to the FILE structure associated with the *stream*.

*File-name* points to a character string that contains the name of the file to be opened.

*Type* is a character string having one of the following values:

"r"	open for reading
"w"	truncate or create for writing
"a"	append; open for writing at end of file, or create for writing
"r+"	open for update (reading and writing)
"w+"	truncate or create for update
"a+"	append; open or create for update at end-of-file

*Freopen* substitutes the named file in place of the open *stream*. The original *stream* is closed, regardless of whether the open ultimately succeeds. *Freopen* returns a pointer to the FILE structure associated with *stream*.

*Freopen* is typically used to attach the preopened *streams* associated with *stdin*, *stdout* and *stderr* to other files.

*Fdopen* associates a *stream* with a file descriptor. File descriptors are obtained from *open*, *dup*, *creat*, or *pipe(2)*, which open files but do not return pointers to a FILE structure *stream*. Streams are necessary input for many of the Section 3S library routines. The *type* of *stream* must agree with the mode of the open file.

When a file is opened for update, both input and output may be done on the resulting *stream*. However, output may not be directly followed by input without an intervening *fseek* or *rewind*, and input may not be directly followed by output without an intervening *fseek*, *rewind*, or an input operation which encounters end-of-file.

When a file is opened for append (i.e., when *type* is "a" or "a+"), it is impossible to overwrite information already in the file. *Fseek* may be used to reposition the file pointer to any position in the file, but when output is written to the file, the current file pointer is disregarded. All output is written at the end of the file and causes the file pointer to be repositioned at the end of the output. If two separate processes open the same file for append, each process may write freely to the file without fear of destroying output being written by the other. The output from the two processes will be intermixed in the file in the order in which it is written.

**SEE ALSO**

`creat(2)`, `dup(2)`, `open(2)`, `pipe(2)`, `fclose(3S)`, `fseek(3S)`.

**DIAGNOSTICS**

*Fopen* and *freopen* return a NULL pointer on failure.

## FREAD(3S)

### NAME

fread, fwrite — binary input/output

### SYNOPSIS

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

```
int fread (ptr, size, nitems, stream)
```

```
char *ptr;
```

```
int size, nitems;
```

```
FILE *stream;
```

```
int fwrite (ptr, size, nitems, stream)
```

```
char *ptr;
```

```
int size, nitems;
```

```
FILE *stream;
```

### DESCRIPTION

*Fread* copies, into an array pointed to by *ptr*, *nitems* items of data from the named input *stream*, where an item of data is a sequence of bytes (not necessarily terminated by a null byte) of length *size*. *Fread* stops appending bytes if an end-of-file or error condition is encountered while reading *stream*, or if *nitems* items have been read. *Fread* leaves the file pointer in *stream*, if defined, pointing to the byte following the last byte read if there is one. *Fread* does not change the contents of *stream*.

*Fwrite* appends at most *nitems* items of data from the array pointed to by *ptr* to the named output *stream*. *Fwrite* stops appending when it has appended *nitems* items of data or if an error condition is encountered on *stream*. *Fwrite* does not change the contents of the array pointed to by *ptr*.

The argument *size* is typically *sizeof(\*ptr)* where the pseudo-function *sizeof* specifies the length of an item pointed to by *ptr*. If *ptr* points to a data type other than *char* it should be cast into a pointer to *char*.

### SEE ALSO

read(2), write(2), fopen(3S), getc(3S), gets(3S), printf(3S), putc(3S), puts(3S), scanf(3S).

### DIAGNOSTICS

*Fread* and *fwrite* return the number of items read or written. If *size* or *nitems* is non-positive, no characters are read or written and 0 is returned by both *fread* and *fwrite*.

### BUGS

On the PDP-11 and iAPX286, the number of bytes transferred is the product of *size* and *nitems*, modulo 65536.

**NAME**

*frexp*, *ldexp*, *modf* — manipulate parts of floating-point numbers

**SYNOPSIS**

**double *frexp* (value, *eptr*)**

**double value;**

**int \**eptr*;**

**double *ldexp* (value, *exp*)**

**double value;**

**int *exp*;**

**double *modf* (value, *iptr*)**

**double value, \**iptr*;**

**DESCRIPTION**

Every non-zero number can be written uniquely as  $x \cdot 2^n$ , where the "mantissa" (fraction)  $x$  is in the range  $0.5 \leq |x| < 1.0$ , and the "exponent"  $n$  is an integer. *Frexp* returns the mantissa of a double *value*, and stores the exponent indirectly in the location pointed to by *eptr*. If *value* is zero, both results returned by *frexp* are zero.

*Ldexp* returns the quantity  $value \cdot 2^{exp}$ .

*Modf* returns the signed fractional part of *value* and stores the integral part indirectly in the location pointed to by *iptr*.

**DIAGNOSTICS**

If *ldexp* would cause overflow,  $\pm$ HUGE is returned (according to the sign of *value*), and *errno* is set to ERANGE.

If *ldexp* would cause underflow, zero is returned and *errno* is set to ERANGE.

S-3

## FSEEK(3S)

### NAME

*fseek*, *rewind*, *ftell* – reposition a file pointer in a stream

### SYNOPSIS

```
#include <stdio.h>

int fseek (stream, offset, ptrname)
FILE *stream;
long offset;
int ptrname;

void rewind (stream)
FILE *stream;

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

### DESCRIPTION

*Fseek* sets the position of the next input or output operation on the *stream*. The new position is at the signed distance *offset* bytes from the beginning, from the current position, or from the end of the file, according as *ptrname* has the value 0, 1, or 2.

*Rewind(stream)* is equivalent to *fseek(stream, 0L, 0)*, except that no value is returned.

*Fseek* and *rewind* undo any effects of *ungetc(3S)*.

After *fseek* or *rewind*, the next operation on a file opened for update may be either input or output.

*Ftell* returns the offset of the current byte relative to the beginning of the file associated with the named *stream*.

### SEE ALSO

*lseek(2)*, *fopen(3S)*, *popen(3S)*, *ungetc(3S)*.

### DIAGNOSTICS

*Fseek* returns non-zero for improper seeks, otherwise zero. An improper seek can be, for example, an *fseek* done on a file that has not been opened via *fopen*; in particular, *fseek* may not be used on a terminal, or on a file opened via *popen(3S)*.

### WARNING

Although on the UNIX system an offset returned by *ftell* is measured in bytes, and it is permissible to seek to positions relative to that offset, portability to non-UNIX systems requires that an offset be used by *fseek* directly. Arithmetic may not meaningfully be performed on such an offset, which is not necessarily measured in bytes.

## NAME

*ftw* — walk a file tree

## SYNOPSIS

```
#include <ftw.h>

int ftw (path, fn, depth)
char *path;
int (*fn) ( );
int depth;
```

## DESCRIPTION

*Ftw* recursively descends the directory hierarchy rooted in *path*. For each object in the hierarchy, *ftw* calls *fn*, passing it a pointer to a null-terminated character string containing the name of the object, a pointer to a *stat* structure [see *stat(2)*] containing information about the object, and an integer. Possible values of the integer, defined in the *<ftw.h>* header file, are FTW\_F for a file, FTW\_D for a directory, FTW\_DNR for a directory that cannot be read, and FTW\_NS for an object for which *stat* could not successfully be executed. If the integer is FTW\_DNR, descendants of that directory will not be processed. If the integer is FTW\_NS, the *stat* structure will contain garbage. An example of an object that would cause FTW\_NS to be passed to *fn* would be a file in a directory with read but without execute (search) permission.

*Ftw* visits a directory before visiting any of its descendants.

The tree traversal continues until the tree is exhausted, an invocation of *fn* returns a nonzero value, or some error is detected within *ftw* (such as an I/O error). If the tree is exhausted, *ftw* returns zero. If *fn* returns a nonzero value, *ftw* stops its tree traversal and returns whatever value was returned by *fn*. If *ftw* detects an error, it returns -1, and sets the error type in *errno*.

*Ftw* uses one file descriptor for each level in the tree. The *depth* argument limits the number of file descriptors so used. If *depth* is zero or negative, the effect is the same as if it were 1. *Depth* must not be greater than the number of file descriptors currently available for use. *Ftw* will run more quickly if *depth* is at least as large as the number of levels in the tree.

## SEE ALSO

*stat(2)*, *malloc(3C)*.

## BUGS

Because *ftw* is recursive, it is possible for it to terminate with a memory fault when applied to very deep file structures.

It could be made to run faster and use less storage on deep structures at the cost of considerable complexity.

*Ftw* uses *malloc(3C)* to allocate dynamic storage during its operation. If *ftw* is forcibly terminated, such as by *longjmp* being executed by *fn* or an interrupt routine, *ftw* will not have a chance to free that storage, so it will remain permanently allocated. A safe way to handle interrupts is to store the fact that an interrupt has occurred, and arrange to have *fn* return a nonzero value at its next invocation.

## GETC(3S)

### NAME

*getc*, *getchar*, *fgetc*, *getw* — get character or word from a stream

### SYNOPSIS

```
#include <stdio.h>
int getc (stream)
FILE *stream;
int getchar ()
int fgetc (stream)
FILE *stream;
int getw (stream)
FILE *stream;
```

### DESCRIPTION

*Getc* returns the next character (i.e., byte) from the named input *stream*, as an integer. It also moves the file pointer, if defined, ahead one character in *stream*. *Getchar* is defined as *getc(stdin)*. *Getc* and *getchar* are macros.

*Fgetc* behaves like *getc*, but is a function rather than a macro. *Fgetc* runs more slowly than *getc*, but it takes less space per invocation and its name can be passed as an argument to a function.

*Getw* returns the next word (i.e., integer) from the named input *stream*. *Getw* increments the associated file pointer, if defined, to point to the next word. The size of a word is the size of an integer and varies from machine to machine. *Getw* assumes no special alignment in the file.

### SEE ALSO

*fclose(3S)*, *feof(3S)*, *fopen(3S)*, *fread(3S)*, *gets(3S)*, *putc(3S)*, *scanf(3S)*.

### DIAGNOSTICS

These functions return the constant EOF at end-of-file or upon an error. Because EOF is a valid integer, *feof(3S)* should be used to detect *getw* errors.

### WARNING

If the integer value returned by *getc*, *getchar*, or *fgetc* is stored into a character variable and then compared against the integer constant EOF, the comparison may never succeed, because sign-extension of a character on widening to integer is machine-dependent.

### BUGS

Because it is implemented as a macro, *getc* treats incorrectly a *stream* argument with side effects. In particular, *getc(\*f++)* does not work sensibly. *Fgetc* should be used instead.

Because of possible differences in word length and byte ordering, files written using *putw* are machine-dependent, and may not be read using *getw* on a different processor.



**NAME**

`getcwd` — get path name of current working directory

**SYNOPSIS**

```
char *getcwd (buf, size)
char *buf;
int size;
```

**DESCRIPTION**

*Getcwd* returns a pointer to the current directory path name. The value of *size* must be at least two greater than the length of the path name to be returned.

If *buf* is a NULL pointer, *getcwd* will obtain *size* bytes of space using *malloc*(3C). In this case, the pointer returned by *getcwd* may be used as the argument in a subsequent call to *free*.

The function is implemented by using *popen*(3S) to pipe the output of the *pwd*(1) command into the specified string space.

**EXAMPLE**

```
char *cwd, *getcwd();
.
.
if ((cwd = getcwd((char *)NULL, 64)) == NULL) {
    perror("pwd");
    exit(1);
}
printf("%s\n", cwd);
```

**SEE ALSO**

*malloc*(3C), *popen*(3S)  
*pwd*(1) in the Runtime System manual. *uul*.

**DIAGNOSTICS**

Returns NULL with *errno* set if *size* is not large enough, or if an error occurs in a lower-level function.

## GETENV(3C)

### NAME

getenv — return value for environment name

### SYNOPSIS

```
char *getenv (name)
char *name;
```

### DESCRIPTION

*Getenv* searches the environment list [see *environ(5)*] for a string of the form *name=value*, and returns a pointer to the *value* in the current environment if such a string is present, otherwise a NULL pointer.

### SEE ALSO

exec(2), putenv(3C), environ(5).

## NAME

*getgrent*, *getgrgid*, *getgrnam*, *setgrent*, *endgrent*, *fgetgrent* - get group file entry

## SYNOPSIS

```
#include <grp.h>

struct group *getgrent ( )
struct group *getgrgid (gid)
int gid;
struct group *getgrnam (name)
char *name;
void setgrent ( )
void endgrent ( )
struct group *fgetgrent (f)
FILE *f;
```

## DESCRIPTION

*Getgrent*, *getgrgid* and *getgrnam* each return pointers to an object with the following structure containing the broken-out fields of a line in the */etc/group* file. Each line contains a "group" structure, defined in the *<grp.h>* header file.

```
struct group {
    char *gr_name; /* the name of the group */
    char *gr_passwd; /* the encrypted group password */
    int gr_gid; /* the numerical group ID */
    char **gr_mem; /* vector of pointers to member names */
};
```

*Getgrent* when first called returns a pointer to the first group structure in the file; thereafter, it returns a pointer to the next group structure in the file; so, successive calls may be used to search the entire file. *Getgrgid* searches from the beginning of the file until a numerical group id matching *gid* is found and returns a pointer to the particular structure in which it was found. *Getgrnam* searches from the beginning of the file until a group name matching *name* is found and returns a pointer to the particular structure in which it was found. If an end-of-file or an error is encountered on reading, these functions return a NULL pointer.

A call to *setgrent* has the effect of rewinding the group file to allow repeated searches. *Endgrent* may be called to close the group file when processing is complete.

*Fgetgrent* returns a pointer to the next group structure in the stream *f*, which matches the format of */etc/group*.

## FILES

*/etc/group*

## SEE ALSO

*getlogin(3C)*, *getpwent(3C)*, *group(4)*.

## DIAGNOSTICS

A NULL pointer is returned on EOF or error.

## WARNING

The above routines use *<stdio.h>*, which causes them to increase the size of programs, not otherwise using standard I/O, more than might be expected.

## BUGS

All information is contained in a static area, so it must be copied if it is to be saved.

## GETLOGIN(3C)

### NAME

getlogin — get login name

### SYNOPSIS

```
char *getlogin ( );
```

### DESCRIPTION

*Getlogin* returns a pointer to the login name as found in */etc/utmp*. It may be used in conjunction with *getpwnam* to locate the correct password file entry when the same user ID is shared by several login names.

If *getlogin* is called within a process that is not attached to a terminal, it returns a NULL pointer. The correct procedure for determining the login name is to call *cuserid*, or to call *getlogin* and if it fails, to call *getpwuid*.

### FILES

*/etc/utmp*

### SEE ALSO

*cuserid(3S)*, *getgrent(3C)*, *getpwent(3C)*, *utmp(4)*.

### DIAGNOSTICS

Returns the NULL pointer if *name* is not found.

### BUGS

The return values point to static data whose content is overwritten by each call.

## NAME

getopt — get option letter from argument vector

## SYNOPSIS

```
int getopt (argc, argv, optstring)
int argc;
char **argv, *optstring;
extern char *optarg;
extern int optind, opterr;
```

## DESCRIPTION

*Getopt* returns the next option letter in *argv* that matches a letter in *optstring*. *Optstring* is a string of recognized option letters; if a letter is followed by a colon, the option is expected to have an argument that may or may not be separated from it by white space. *Optarg* is set to point to the start of the option argument on return from *getopt*.

*Getopt* places in *optind* the *argv* index of the next argument to be processed. Because *optind* is external, it is normally initialized to zero automatically before the first call to *getopt*.

When all options have been processed (i.e., up to the first non-option argument), *getopt* returns EOF. The special option -- may be used to delimit the end of the options; EOF will be returned, and -- will be skipped.

## DIAGNOSTICS

*Getopt* prints an error message on *stderr* and returns a question mark (?) when it encounters an option letter not included in *optstring*. This error message may be disabled by setting *opterr* to a non-zero value.

## EXAMPLE

The following code fragment shows how one might process the arguments for a command that can take the mutually exclusive options a and b, and the options f and o, both of which require arguments:

```
main (argc, argv)
int argc;
char **argv;
{
    int c;
    extern char *optarg;
    extern int optind;
    .
    .
    .
    while ((c = getopt(argc, argv, "abf:o:")) != EOF)
        switch (c) {
            case 'a':
                if (bflag)
                    errflg++;
                else
                    aflag++;
                break;
            case 'b':
                if (aflag)
                    errflg++;
                else
                    bproc( );
                break;
            case 'f':
```

## GETOPT(3C)

```
        ifile = optarg;
        break;
    case 'o':
        ofile = optarg;
        break;
    case '?':
        errflg++;
    }
    if (errflg) {
        fprintf(stderr, "usage: . . . ");
        exit (2);
    }
    for ( ; optind < argc; optind++) {
        if (access(argv[optind], 4)) {
            .
            .
            .
        }
    }
```

### SEE ALSO

getopt(1) in the Runtime System manual.

**NAME**

getpass — read a password

**SYNOPSIS**

```
char *getpass (prompt)
char *prompt;
```

**DESCRIPTION**

*Getpass* reads up to a new-line or EOF from the file */dev/tty*, after prompting on the standard error output with the null-terminated string *prompt* and disabling echoing. A pointer is returned to a null-terminated string of at most 8 characters. If */dev/tty* cannot be opened, a NULL pointer is returned. An interrupt will terminate input and send an interrupt signal to the calling program before returning.

**FILES**

*/dev/tty*

**SEE ALSO**

crypt(3C).

**WARNING**

The above routine uses `<stdio.h>`, which causes it to increase the size of programs not otherwise using standard I/O, more than might be expected.

**BUGS**

The return value points to static data whose content is overwritten by each call.

## GETPW(3C)

### NAME

getpw — get name from UID

### SYNOPSIS

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

### DESCRIPTION

*Getpw* searches the password file for a user id number that equals *uid*, copies the line of the password file in which *uid* was found into the array pointed to by *buf*, and returns 0. *Getpw* returns non-zero if *uid* cannot be found.

This routine is included only for compatibility with prior systems and should not be used; see *getpwent*(3C) for routines to use instead.

### FILES

/etc/passwd

### SEE ALSO

getpwent(3C), passwd(4).

### DIAGNOSTICS

*Getpw* returns non-zero on error.

### WARNING

The above routine uses `<stdio.h>`, which causes it to increase, more than might be expected, the size of programs not otherwise using standard I/O.



## NAME

getpwent, getpwuid, getpwnam, setpwent, endpwent, fgetpwent — get password file entry

## SYNOPSIS

```
#include <pwd.h>

struct passwd *getpwent ( )
struct passwd *getpwuid (uid)
int uid;

struct passwd *getpwnam (name)
char *name;

void setpwent ( )
void endpwent ( )

struct passwd *fgetpwent (f)
FILE *f;
```

## DESCRIPTION

*Getpwent*, *getpwuid* and *getpwnam* each returns a pointer to an object with the following structure containing the broken-out fields of a line in the */etc/passwd* file. Each line in the file contains a “passwd” structure, declared in the *<pwd.h>* header file:

```
struct passwd {
    char    *pw_name;
    char    *pw_passwd;
    int     pw_uid;
    int     pw_gid;
    char    *pw_age;
    char    *pw_comment;
    char    *pw_gecos;
    char    *pw_dir;
    char    *pw_shell;
};
```

This structure is declared in *<pwd.h>* so it is not necessary to redeclare it.

The *pw\_comment* field is unused; the others have meanings described in *passwd(4)*.

*Getpwent* when first called returns a pointer to the first passwd structure in the file; thereafter, it returns a pointer to the next passwd structure in the file; so successive calls can be used to search the entire file. *Getpwuid* searches from the beginning of the file until a numerical user id matching *uid* is found and returns a pointer to the particular structure in which it was found. *Getpwnam* searches from the beginning of the file until a login name matching *name* is found, and returns a pointer to the particular structure in which it was found. If an end-of-file or an error is encountered on reading, these functions return a NULL pointer.

A call to *setpwent* has the effect of rewinding the password file to allow repeated searches. *Endpwent* may be called to close the password file when processing is complete.

*Fgetpwent* returns a pointer to the next passwd structure in the stream *f*, which matches the format of */etc/passwd*.

## FILES

*/etc/passwd*

## GETPWENT(3C)

### SEE ALSO

getlogin(3C), getgrent(3C), passwd(4).

### DIAGNOSTICS

A NULL pointer is returned on EOF or error.

### WARNING

The above routines use `<stdio.h>`, which causes them to increase the size of programs, not otherwise using standard I/O, more than might be expected.

### BUGS

All information is contained in a static area, so it must be copied if it is to be saved.

## NAME

gets, fgets — get a string from a stream

## SYNOPSIS

```
#include <stdio.h>

char *gets (s)
char *s;

char *fgets (s, n, stream)
char *s;
int n;
FILE *stream;
```

## DESCRIPTION

*Gets* reads characters from the standard input stream, *stdin*, into the array pointed to by *s*, until a new-line character is read or an end-of-file condition is encountered. The new-line character is discarded and the string is terminated with a null character.

*Fgets* reads characters from the *stream* into the array pointed to by *s*, until *n*-1 characters are read, or a new-line character is read and transferred to *s*, or an end-of-file condition is encountered. The string is then terminated with a null character.

## SEE ALSO

ferror(3S), fopen(3S), fread(3S),getc(3S), scanf(3S).

## DIAGNOSTICS

If end-of-file is encountered and no characters have been read, no characters are transferred to *s* and a NULL pointer is returned. If a read error occurs, such as trying to use these functions on a file that has not been opened for reading, a NULL pointer is returned. Otherwise *s* is returned.

# GETUT(3C)

## NAME

getutent, getutid, getutline, pututline, setutent, endutent, utmpname - access utmp file entry

## SYNOPSIS

```
#include <utmp.h>

struct utmp *getutent ( )
struct utmp *getutid (id)
struct utmp *id;
struct utmp *getutline (line)
struct utmp *line;
void pututline (utmp)
struct utmp *utmp;
void setutent ( )
void endutent ( )
void utmpname (file)
char *file;
```

## DESCRIPTION

*Getutent*, *getutid* and *getutline* each return a pointer to a structure of the following type:

```
struct utmp {
    char    ut_user[8];        /* User login name */
    char    ut_id[4];         /* /etc/inittab id
                             * (usually line #) */
    char    ut_line[12];      /* device name (console,
                             * lnxx) */
    short   ut_pid;           /* process id */
    short   ut_type;          /* type of entry */
    struct  exit_status {
        short e_termination; /* Process termination status */
        short e_exit;        /* Process exit status */
    } ut_exit;                /* The exit status of a process
                             * marked as DEAD_PROCESS. */
    time_t  ut_time;          /* time entry was made */
};
```

*Getutent* reads in the next entry from a *utmp*-like file. If the file is not already open, it opens it. If it reaches the end of the file, it fails.

*Getutid* searches forward from the current point in the *utmp* file until it finds an entry with a *ut\_type* matching *id* -> *ut\_type* if the type specified is RUN\_LVL, BOOT\_TIME, OLD\_TIME or NEW\_TIME. If the type specified in *id* is INIT\_PROCESS, LOGIN\_PROCESS, USER\_PROCESS or DEAD\_PROCESS, then *getutid* will return a pointer to the first entry whose type is one of these four and whose *ut\_id* field matches *id* -> *ut\_id*. If the end-of-file is reached without a match, it fails.

*Getutline* searches forward from the current point in the *utmp* file until it finds an entry of the type LOGIN\_PROCESS or USER\_PROCESS which also has a *ut\_line* string matching the *line* -> *ut\_line* string. If the end-of-file is reached without a match, it fails.

*Pututline* writes out the supplied *utmp* structure into the *utmp* file. It uses *getutid* to search forward for the proper place if it finds that it is not already at the proper place. It is expected that normally the user of *pututline* will have searched for the proper entry using one of the *getut* routines. If so, *pututline*

will not search. If *pututline* does not find a matching slot for the new entry, it will add a new entry to the end of the file.

*Setutent* resets the input stream to the beginning of the file. This should be done before each search for a new entry if it is desired that the entire file be examined.

*Endutent* closes the currently open file.

*Utmpname* allows the user to change the name of the file examined, from */etc/utmp* to any other file. It is most often expected that this other file will be */etc/wtmp*. If the file does not exist, this will not be apparent until the first attempt to reference the file is made. *Utmpname* does not open the file. It just closes the old file if it is currently open and saves the new file name.

#### FILES

*/etc/utmp*  
*/etc/wtmp*

#### SEE ALSO

*ttyslot(3C)*, *utmp(4)*.

#### DIAGNOSTICS

A NULL pointer is returned upon failure to read, whether for permissions or having reached the end of file, or upon failure to write.

#### COMMENTS

The most current entry is saved in a static structure. Multiple accesses require that it be copied before further accesses are made. Each call to either *getutid* or *getutline* sees the routine examine the static structure before performing more I/O. If the contents of the static structure match what it is searching for, it looks no further. For this reason to use *getutline* to search for multiple occurrences, it would be necessary to zero out the static after each success, or *getutline* would just return the same pointer over and over again. There is one exception to the rule about removing the structure before further reads are done. The implicit read done by *pututline* (if it finds that it is not already at the correct place in the file) will not hurt the contents of the static structure returned by the *getutent*, *getutid* or *getutline* routines, if the user has just modified those contents and passed the pointer back to *pututline*.

These routines use buffered standard I/O for input, but *pututline* uses an unbuffered non-standard write to avoid race conditions between processes trying to modify the *utmp* and *wtmp* files.

# HSEARCH(3C)

## NAME

`hsearch`, `hcreate`, `hdestroy` — manage hash search tables

## SYNOPSIS

```
#include <search.h>

ENTRY *hsearch (item, action)
ENTRY item;
ACTION action;

int hcreate (nel)
unsigned nel;

void hdestroy ( )
```

## DESCRIPTION

*Hsearch* is a hash-table search routine generalized from Knuth (6.4) Algorithm D. It returns a pointer into a hash table indicating the location at which an entry can be found. *Item* is a structure of type `ENTRY` (defined in the `<search.h>` header file) containing two pointers: *item.key* points to the comparison key, and *item.data* points to any other data to be associated with that key. (Pointers to types other than character should be cast to pointer-to-character.) *Action* is a member of an enumeration type `ACTION` indicating the disposition of the entry if it cannot be found in the table. `ENTER` indicates that the item should be inserted in the table at an appropriate point. `FIND` indicates that no entry should be made. Unsuccessful resolution is indicated by the return of a `NULL` pointer.

*Hcreate* allocates sufficient space for the table, and must be called before *hsearch* is used. *Nel* is an estimate of the maximum number of entries that the table will contain. This number may be adjusted upward by the algorithm in order to obtain certain mathematically favorable circumstances.

*Hdestroy* destroys the search table, and may be followed by another call to *hcreate*.

## NOTES

*Hsearch* uses *open addressing* with a *multiplicative* hash function. However, its source code has many other options available which the user may select by compiling the *hsearch* source with the following symbols defined to the preprocessor:

- DIV** Use the *remainder modulo table size* as the hash function instead of the multiplicative algorithm.
- USCR** Use a User-Supplied Comparison Routine for ascertaining table membership. The routine should be named *hcompare* and should behave in a manner similar to *strcmp* [see *string(3C)*].
- CHAINED** Use a linked list to resolve collisions. If this option is selected, the following other options become available.
  - START** Place new entries at the beginning of the linked list (default is at the end).
  - SORTUP** Keep the linked list sorted by key in ascending order.
  - SORTDOWN** Keep the linked list sorted by key in descending order.

Additionally, there are preprocessor flags for obtaining debugging printout (`-DDEBUG`) and for including a test driver in the calling routine (`-DDRIVER`). The source code should be consulted for further details.

## EXAMPLE

The following example will read in strings followed by two numbers and store them in a hash table, discarding duplicates. It will then read in strings and find the matching entry in the hash table and print it out.

```

#include <stdio.h>
#include <search.h>

struct info {          /* this is the info stored in the table */
    int age, room; /* other than the key. */
};
#define NUM_EMPL      5000 /* # of elements in search table */

main()
[
    /* space to store strings */
    char string_space[NUM_EMPL*20];
    /* space to store employee info */
    struct info info_space[NUM_EMPL];
    /* next avail space in string_space */
    char *str_ptr = string_space;
    /* next avail space in info_space */
    struct info *info_ptr = info_space;
    ENTRY item, *found_item, *hsearch( );
    /* name to look for in table */
    char name_to_find[30];
    int i = 0;

    /* create table */
    (void) hcreate(NUM_EMPL);
    while (scanf("%s%d%d", str_ptr, &info_ptr->age,
                &info_ptr->room) != EOF && i++ < NUM_EMPL) {
        /* put info in structure, and structure in item */
        item.key = str_ptr;
        item.data = (char *)info_ptr;
        str_ptr += strlen(str_ptr) + 1;
        info_ptr++;
        /* put item into table */
        (void) hsearch(item, ENTER);
    }

    /* access table */
    item.key = name_to_find;
    while (scanf("%s", item.key) != EOF) {
        if ((found_item = hsearch(item, FIND)) != NULL) {
            /* if item is in the table */
            (void)printf("found %s, age = %d, room = %d\n",
                found_item->key,
                ((struct info *)found_item->data)->age,
                ((struct info *)found_item->data)->room);
        } else {
            (void)printf("no such employee %s\n",
                name_to_find)
        }
    }
}
]

```

## HSEARCH(3C)

### SEE ALSO

bsearch(3C), lsearch(3C), malloc(3C), malloc(3X), string(3C), tsearch(3C).

### DIAGNOSTICS

*Hsearch* returns a NULL pointer if either the action is **FIND** and the item could not be found or the action is **ENTER** and the table is full.

*Hcreate* returns zero if it cannot allocate sufficient space for the table.

### WARNING

*Hsearch* and *hcreate* use *malloc(3C)* to allocate space.

### BUGS

Only one hash search table may be active at any given time.



table of active locks. If this table is full, an [EDEADLK] error is returned and the requested section is not released.

A potential for deadlock occurs if a process controlling a locked resource is put to sleep by accessing another process's locked resource. Thus calls to *lock* or *fcntl* scan for a deadlock prior to sleeping on a locked resource. An error return is made if sleeping on the locked resource would cause a deadlock.

Sleeping on a resource is interrupted with any signal. The *alarm(2)* command may be used to provide a timeout facility in applications which require this facility.

The *lockf* utility will fail if one or more of the following are true:

[EBADF]

*Fildes* is not a valid open descriptor.

[EACCESS]

*Cmd* is F\_TLOCK or F\_TEST and the section is already locked by another process.

[EDEADLK]

*Cmd* is F\_LOCK or F\_TLOCK and a deadlock would occur. Also the *cmd* is either of the above or F\_ULOCK and the number of entries in the lock table would exceed the number allocated on the system.

#### SEE ALSO

*alarm(2)*, *close(2)*, *creat(2)*, *fcntl(2)*, *intro(2)*, *open(2)*, *read(2)*, *write(2)*.

#### RETURN VALUE

Upon successful completion, a value of 0 is returned. Otherwise, a value of -1 is returned and *errno* is set to indicate the error.

#### CAVEATS

Unexpected results may occur in processes that do buffering in the user address space. The process may later read/write data which is/was locked. The standard I/O package is the most common source of unexpected buffering.

## LSEARCH(3C)

### NAME

`lsearch`, `lfind` - linear search and update

### SYNOPSIS

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

char *lsearch ((char *)key, (char *)base, nel, sizeof(*key),
compar)
unsigned *nel;
int (*compar)( );

char *lfind ((char *)key, (char *)base, nel, sizeof(*key), com-
par)
unsigned *nel;
int (*compar)( );
```

### DESCRIPTION

*Lsearch* is a linear search routine generalized from Knuth (6.1) Algorithm S. It returns a pointer into a table indicating where a datum may be found. If the datum does not occur, it is added at the end of the table. **Key** points to the datum to be sought in the table. **Base** points to the first element in the table. **Nel** points to an integer containing the current number of elements in the table. The variable pointed to by **nel** is incremented if the datum is added to the table. **Compar** is the name of the comparison function which the user must supply (*strcmp*, for example). It is called with two arguments that point to the elements being compared. The function must return zero if the elements are equal and non-zero otherwise.

*Lfind* is the same as *lsearch* except that if the datum is not found, it is not added to the table. Instead, a NULL pointer is returned.

### NOTES

The pointers to the key and the element at the base of the table should be of type pointer-to-element, and cast to type pointer-to-character. The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

Although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

### EXAMPLE

This fragment will read in  $\leq$  TABSIZE strings of length  $\leq$  ELSIZE and store them in a table, eliminating duplicates.

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

#define TABSIZE 50
#define ELSIZE 120

char line[ELSIZE], tab[TABSIZE][ELSIZE], *lsearch( );
unsigned nel = 0;
int strcmp( );
...
while (fgets(line, ELSIZE, stdin) != NULL &&
    nel < TABSIZE)
    (void) lsearch(line, (char *)tab, &nel,
        ELSIZE, strcmp);
...
```

**NAME**

*l3tol*, *ltol3* - convert between three-byte integers and long integers

**SYNOPSIS**

```
void l3tol (lp, cp, n)
```

```
long *lp;
```

```
char *cp;
```

```
int n;
```

```
void ltol3 (cp, lp, n)
```

```
char *cp;
```

```
long *lp;
```

```
int n;
```

**DESCRIPTION**

*L3tol* converts a list of *n* three-byte integers packed into a character string pointed to by *cp* into a list of long integers pointed to by *lp*.

*Ltol3* performs the reverse conversion from long integers (*lp*) to three-byte integers (*cp*).

These functions are useful for file-system maintenance where the block numbers are three bytes long.

**SEE ALSO**

*fs(4)*.

**BUGS**

Because of possible differences in byte ordering, the numerical values of the long integers are machine-dependent.

## LOCKF(3C)

### NAME

lockf - record locking on files

### SYNOPSIS

```
# include <unistd.h>
```

```
lockf (fildes, function, size) long size; int fildes, function;
```

### DESCRIPTION

The *lockf* command will allow sections of a file to be locked (advisory write locks). (Mandatory or enforcement mode record locks are not currently available.) Locking calls from other processes which attempt to lock the locked file section will either return an error value or be put to sleep until the resource becomes unlocked. All the locks for a process are removed when the process terminates. See *fcntl(2)* for more information about record locking.

*Fildes* is an open file descriptor. The file descriptor must have O\_WRONLY or O\_RDWR permission in order to establish lock with this function call.

*Function* is a control value which specifies the action to be taken. The permissible values for *function* are defined in <unistd.h> as follows:

```
#define F_ULOCK 0 /* Unlock a previously locked section */
#define F_LOCK 1 /* Lock a section for exclusive use */
#define F_TLOCK 2 /* Test and lock a section for exclusive use */
#define F_TEST 3 /* Test section for other processes locks */
```

All other values of *function* are reserved for future extensions and will result in an error return if not implemented.

F\_TEST is used to detect if a lock by another process is present on the specified section. F\_LOCK and F\_TLOCK both lock a section of a file if the section is available. F\_UNLOCK removes locks from a section of the file.

*Size* is the number of contiguous bytes to be locked or unlocked. The resource to be locked starts at the current offset in the file and extends forward for a positive size and backward for a negative size (the preceding bytes up to but not including the current offset). If *size* is zero, the section from the current offset through the largest file offset is locked (i.e., from the current offset through the present or any future end-of-file). An area need not be allocated to the file in order to be locked, as such locks may exist past the end-of-file.

The sections locked with F\_LOCK or F\_TLOCK may, in whole or in part, contain or be contained by a previously locked section for the same process. When this occurs, or if adjacent sections occur, the sections are combined into a single section. If the request requires that a new element be added to the table of active locks and this table is already full, an error is returned, and the new section is not locked.

F\_LOCK and F\_TLOCK requests differ only by the action taken if the resource is not available. F\_LOCK will cause the calling process to sleep until the resource is available. F\_TLOCK will cause the function to return a -1 and set *errno* to [EACCESS] error if the section is already locked by another process.

F\_ULOCK requests may, in whole or in part, release one or more locked sections controlled by the process. When sections are not fully released, the remaining sections are still locked by the process. Releasing the center section of a locked section requires an additional element in the

**SEE ALSO**

bsearch(3C), hsearch(3C), tsearch(3C).

**DIAGNOSTICS**

If the searched for datum is found, both *lsearch* and *lfind* return a pointer to it. Otherwise, *lfind* returns NULL and *lsearch* returns a pointer to the newly added element.

**BUGS**

Undefined results can occur if there is not enough room in the table to add a new item.

# MALLOC(3C)

## NAME

`malloc`, `free`, `realloc`, `calloc` — main memory allocator

## SYNOPSIS

```
char *malloc (size)
    unsigned size;
void free (ptr)
char *ptr;
char *realloc (ptr, size)
char *ptr;
    unsigned size;
char *calloc (nelem, elsize)
    unsigned nelem, elsize;
```

## DESCRIPTION

*Malloc* and *free* provide a simple general-purpose memory allocation package. *Malloc* returns a pointer to a block of at least *size* bytes suitably aligned for any use.

The argument to *free* is a pointer to a block previously allocated by *malloc*; after *free* is performed this space is made available for further allocation, but its contents are left undisturbed.

Undefined results will occur if the space assigned by *malloc* is overrun or if some random number is handed to *free*.

*Malloc* allocates the first big enough contiguous reach of free space found in a circular search from the last block allocated or freed, coalescing adjacent free blocks as it searches. It calls *sbrk* [see *brk(2)*] to get more memory from the system when there is no suitable space already free.

*Realloc* changes the size of the block pointed to by *ptr* to *size* bytes and returns a pointer to the (possibly moved) block. The contents will be unchanged up to the lesser of the new and old sizes. If no free block of *size* bytes is available in the storage arena, then *realloc* will ask *malloc* to enlarge the arena by *size* bytes and will then move the data to the new space.

*Realloc* also works if *ptr* points to a block freed since the last call of *malloc*, *realloc*, or *calloc*; thus sequences of *free*, *malloc* and *realloc* can exploit the search strategy of *malloc* to do storage compaction.

*Calloc* allocates space for an array of *nelem* elements of size *elsize*. The space is initialized to zeros.

Each of the allocation routines returns a pointer to space suitably aligned (after possible pointer coercion) for storage of any type of object.

## SEE ALSO

*brk(2)*, *malloc(3X)*.

## DIAGNOSTICS

*Malloc*, *realloc* and *calloc* return a NULL pointer if there is no available memory or if the arena has been detectably corrupted by storing outside the bounds of a block. When this happens the block pointed to by *ptr* may be destroyed.

## NOTE

Search time increases when many objects have been allocated; that is, if a program allocates but never frees, then each successive allocation takes longer. For an alternate, more flexible implementation, see *malloc(3X)*.

## NAME

memccpy, memchr, memcmp, memcpy, memset — memory operations

## SYNOPSIS

```
#include <memory.h>

char *memccpy (s1, s2, c, n)
char *s1, *s2;
int c, n;

char *memchr (s, c, n)
char *s;
int c, n;

int memcmp (s1, s2, n)
char *s1, *s2;
int n;

char *memcpy (s1, s2, n)
char *s1, *s2;
int n;

char *memset (s, c, n)
char *s;
int c, n;
```

## DESCRIPTION

These functions operate as efficiently as possible on memory areas (arrays of characters bounded by a count, not terminated by a null character). They do not check for the overflow of any receiving memory area.

*Memccpy* copies characters from memory area *s2* into *s1*, stopping after the first occurrence of character *c* has been copied, or after *n* characters have been copied, whichever comes first. It returns a pointer to the character after the copy of *c* in *s1*, or a NULL pointer if *c* was not found in the first *n* characters of *s2*.

*Memchr* returns a pointer to the first occurrence of character *c* in the first *n* characters of memory area *s*, or a NULL pointer if *c* does not occur.

*Memcmp* compares its arguments, looking at the first *n* characters only, and returns an integer less than, equal to, or greater than 0, according as *s1* is lexicographically less than, equal to, or greater than *s2*.

*Memcpy* copies *n* characters from memory area *s2* to *s1*. It returns *s1*.

*Memset* sets the first *n* characters in memory area *s* to the value of character *c*. It returns *s*.

## NOTE

For user convenience, all these functions are declared in the optional *<memory.h>* header file.

## BUGS

*Memcmp* uses native character comparison, which is signed on PDP-11s and VAX-11s, unsigned on other machines. Thus the sign of the value returned when one of the characters has its high-order bit set is implementation-dependent.

Character movement is performed differently in different implementations. Thus overlapping moves may yield surprises.

## MKTEMP(3C)

### NAME

mktemp — make a unique file name

### SYNOPSIS

char \*mktemp (template)

char \*template;

### DESCRIPTION

*Mktemp* replaces the contents of the string pointed to by *template* by a unique file name, and returns the address of *template*. The string in *template* should look like a file name with six trailing Xs; *mktemp* will replace the Xs with a letter and the current process ID. The letter will be chosen so that the resulting name does not duplicate an existing file.

### SEE ALSO

getpid(2), tmpfile(3S), tmpnam(3S).

### BUGS

It is possible to run out of letters.



## NAME

monitor — prepare execution profile

## SYNOPSIS

```
#include <mon.h>

void monitor (lowpc, highpc, buffer, bufsize, nfunc)
int (*lowpc)(), (*highpc)();
WORD *buffer;
int bufsize, nfunc;
```

## DESCRIPTION

An executable program created by `cc -p` automatically includes calls for *monitor* with default parameters; *monitor* needn't be called explicitly except to gain fine control over profiling.

*Monitor* is an interface to *profil(2)*. *Lowpc* and *highpc* are the addresses of two functions; *buffer* is the address of a (user-supplied) array of *bufsize* WORDs (defined in the *<mon.h>* header file). *Monitor* arranges to record a histogram of periodically sampled values of the program counter, and of counts of calls of certain functions, in the buffer. The lowest address sampled is that of *lowpc* and the highest is just below *highpc*. *Lowpc* may not equal 0 for this use of *monitor*. At most *nfunc* call counts can be kept; only calls of functions compiled with the profiling option `-p` of *cc(1)* are recorded. (The C Library and Math Library supplied when `cc -p` is used also have call counts recorded.)

For the results to be significant, especially where there are small, heavily used routines, it is suggested that the buffer be no more than a few times smaller than the range of locations sampled.

To profile the entire program, it is sufficient to use

```
extern void etext0;

...
monitor ((void (*)0)2, etext, buf, bufsize, nfunc);
```

*Etect* lies just above all the program text; see *end(3C)*.

To stop execution monitoring and write the results on the file *mon.out*, use

```
monitor ((void (*)0)0, (void (*)0)0, (word *)0, 0, 0);
```

*Prof(1)* can then be used to examine the results.

## FILES

```
mon.out
/lib/small/libp/libc.a
/lib/small/libp/libm.a
/lib/large/libp/libc.a
/lib/large/libp/libm.a
/lib/huge/libp/libc.a
/lib/huge/libp/libm.a
```

## SEE ALSO

*profil(2)*, *end(3C)*,  
*cc(1)*, *prof(1)* in the Runtime System manual.

S-3

## NLIST(3C)

### NAME

*nlist* — get entries from name list

### SYNOPSIS

```
#include <nlist.h>
int nlist (file-name, nl)
char *file-name;
struct nlist *nl;
```

### DESCRIPTION

*Nlist* examines the name list in the executable file whose name is pointed to by *file-name*, and selectively extracts a list of values and puts them in the array of *nlist* structures pointed to by *nl*. The name list *nl* consists of an array of structures containing names of variables, types and values. The list is terminated with a null name; that is, a null string is in the name position of the structure. Each variable name is looked up in the name list of the file. If the name is found, the type and value of the name are inserted in the next two fields. The type field will be set to 0 unless the file was compiled with the *-g* option. If the name is not found, both entries are set to 0. See *a.out(4)* for a discussion of the symbol table structure.

This function is useful for examining the system name list kept in the file */unix*. In this way programs can obtain system addresses that are up to date.

### NOTES

The *<nlist.h>* header file is automatically included by *<a.out.h>* for compatibility. However, if the only information needed from *<a.out.h>* is for use of *nlist*, then including *<a.out.h>* is discouraged. If *<a.out.h>* is included, the line *"#undef n\_name"* may need to follow it.

### SEE ALSO

*a.out(4)*.

### DIAGNOSTICS

All value entries are set to 0 if the file cannot be read or if it does not contain a valid name list.

*Nlist* returns *-1* upon error; otherwise it returns 0.

**NAME**

`error`, `errno`, `sys_errlist`, `sys_nerr` — system error messages

**SYNOPSIS**

```
void error (s)
char *s;
extern int errno;
extern char *sys_errlist [];
extern int sys_nerr;
```

**DESCRIPTION**

*error* produces a message on the standard error output, describing the last error encountered during a call to a system or library function. The argument string *s* is printed first, then a colon and a blank, then the message and a new-line. To be of most use, the argument string should include the name of the program that incurred the error. The error number is taken from the external variable *errno*, which is set when errors occur but not cleared when non-erroneous calls are made.

To simplify variant formatting of messages, the array of message strings *sys\_errlist* is provided; *errno* can be used as an index in this table to get the message string without the new-line. *sys\_nerr* is the largest message number provided for in the table; it should be checked because new error codes may be added to the system before they are added to the table.

**SEE ALSO**

`intro(2)`.

## POPEN(3S)

### NAME

*popen*, *pclose* — initiate pipe to/from a process

### SYNOPSIS

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

```
FILE *popen (command, type)
```

```
char *command, *type;
```

```
int pclose (stream)
```

```
FILE *stream;
```

### DESCRIPTION

The arguments to *popen* are pointers to null-terminated strings containing, respectively, a shell command line and an I/O mode, either *r* for reading or *w* for writing. *Popen* creates a pipe between the calling program and the command to be executed. The value returned is a stream pointer such that one can write to the standard input of the command, if the I/O mode is *w*, by writing to the file *stream*; and one can read from the standard output of the command, if the I/O mode is *r*, by reading from the file *stream*.

A stream opened by *popen* should be closed by *pclose*, which waits for the associated process to terminate and returns the exit status of the command.

Because open files are shared, a type *r* command may be used as an input filter and a type *w* as an output filter.

### SEE ALSO

*pipe(2)*, *wait(2)*, *fclose(3S)*, *fopen(3S)*, *system(3S)*.

### DIAGNOSTICS

*Popen* returns a NULL pointer if files or processes cannot be created, or if the shell cannot be accessed.

*Pclose* returns -1 if *stream* is not associated with a "*popened*" command.

### BUGS

If the original and "*popened*" processes concurrently read or write a common file, neither should use buffered I/O, because the buffering gets all mixed up. Problems with an output filter may be forestalled by careful buffer flushing, e.g., with *flush*; see *fclose(3S)*.

NAME

printf, fprintf, sprintf — print formatted output

SYNOPSIS

```
#include <stdio.h>

int printf (format [ , arg ] ... )
char *format;

int fprintf (stream, format [ , arg ] ... )
FILE *stream;
char *format;

int sprintf (s, format [ , arg ] ... )
char *s, format;
```

DESCRIPTION

*Printf* places output on the standard output stream *stdout*. *Fprintf* places output on the named output *stream*. *Sprintf* places "output," followed by the null character (\0), in consecutive bytes starting at \*s; it is the user's responsibility to ensure that enough storage is available. Each function returns the number of characters transmitted (not including the \0 in the case of *sprintf*), or a negative value if an output error was encountered.

Each of these functions converts, formats, and prints its *args* under control of the *format*. The *format* is a character string that contains two types of objects: plain characters, which are simply copied to the output stream, and conversion specifications, each of which results in fetching of zero or more *args*. The results are undefined if there are insufficient *args* for the format. If the format is exhausted while *args* remain, the excess *args* are simply ignored.

Each conversion specification is introduced by the character %. After the %, the following appear in sequence:

Zero or more *flags*, which modify the meaning of the conversion specification.

An optional decimal digit string specifying a minimum *field width*. If the converted value has fewer characters than the field width, it will be padded on the left (or right, if the left-adjustment flag '-', described below, has been given) to the field width. If the field width for an s conversion is preceded by a 0, the string is right adjusted with zero-padding on the left.

A *precision* that gives the minimum number of digits to appear for the d, o, u, x, or X conversions, the number of digits to appear after the decimal point for the e and f conversions, the maximum number of significant digits for the g conversion, or the maximum number of characters to be printed from a string in s conversion. The precision takes the form of a period (.) followed by a decimal digit string; a null digit string is treated as zero.

An optional l (ell) specifying that a following d, o, u, x, or X conversion character applies to a long integer *arg*. A l before any other conversion character is ignored.

A character that indicates the type of conversion to be applied.

A field width or precision may be indicated by an asterisk (\*) instead of a digit string. In this case, an integer *arg* supplies the field width or precision. The *arg* that is actually converted is not fetched until the conversion letter is seen, so the *args* specifying field width or precision must appear *before* the *arg* (if any) to be converted.

## PRINTF(3S)

The flag characters and their meanings are:

- The result of the conversion will be left-justified within the field.
- + The result of a signed conversion will always begin with a sign (+ or -).
- blank If the first character of a signed conversion is not a sign, a blank will be prefixed to the result. This implies that if the blank and + flags both appear, the blank flag will be ignored.
- # This flag specifies that the value is to be converted to an "alternate form." For **c**, **d**, **s**, and **u** conversions, the flag has no effect. For **o** conversion, it increases the precision to force the first digit of the result to be a zero. For **x** or **X** conversion, a non-zero result will have **0x** or **0X** prefixed to it. For **e**, **E**, **f**, **g**, and **G** conversions, the result will always contain a decimal point, even if no digits follow the point (normally, a decimal point appears in the result of these conversions only if a digit follows it). For **g** and **G** conversions, trailing zeroes will *not* be removed from the result (which they normally are).

The conversion characters and their meanings are:

- d,o,u,x,X** The integer *arg* is converted to signed decimal, unsigned octal, decimal, or hexadecimal notation (**x** and **X**), respectively; the letters **abcdef** are used for **x** conversion and the letters **ABCDEF** for **X** conversion. The precision specifies the minimum number of digits to appear; if the value being converted can be represented in fewer digits, it will be expanded with leading zeroes. (For compatibility with older versions, padding with leading zeroes may alternatively be specified by prepending a zero to the field width. This does not imply an octal value for the field width.) The default precision is 1. The result of converting a zero value with a precision of zero is a null string.
- f** The float or double *arg* is converted to decimal notation in the style "[ - ]ddd.ddd," where the number of digits after the decimal point is equal to the precision specification. If the precision is missing, six digits are output; if the precision is explicitly 0, no decimal point appears.
- e,E** The float or double *arg* is converted in the style "[ - ]d.ddde±dd," where there is one digit before the decimal point and the number of digits after it is equal to the precision; when the precision is missing, six digits are produced; if the precision is zero, no decimal point appears. The **E** format code will produce a number with **E** instead of **e** introducing the exponent. The exponent always contains at least two digits.
- g,G** The float or double *arg* is printed in style **f** or **e** (or in style **E** in the case of a **G** format code), with the precision specifying the number of significant digits. The style used depends on the value converted: style **e** will be used only if the exponent resulting from the conversion is less than -4 or greater than the precision. Trailing zeroes are removed from the result; a decimal point appears only if it is followed by a digit.
- c** The character *arg* is printed.
- s** The *arg* is taken to be a string (character pointer) and characters from the string are printed until a null character (**\0**) is encountered or the number of characters indicated by the precision specification is reached. If the precision is missing, it is taken to be infinite, so all characters up to the first null character are printed. A **NULL** value for *arg* will yield undefined results.

%           Print a %; no argument is converted.

In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result. Characters generated by *printf* and *fprintf* are printed as if *putc(3S)* had been called.

#### EXAMPLES

To print a date and time in the form "Sunday, July 3, 10:02," where *weekday* and *month* are pointers to null-terminated strings:

```
printf(" %s, %s %d, %d:%.2d", weekday, month, day, hour, min);
```

To print  $\pi$  to 5 decimal places:

```
printf(" pi = %.5f", 4 * atan(1.0));
```

#### SEE ALSO

*ecvt(3C)*, *putc(3S)*, *scanf(3S)*, *stdio(3S)*.

## PUTC(3S)

### NAME

`putc`, `putchar`, `fputc`, `putw` - put character or word on a stream

### SYNOPSIS

```
#include <stdio.h>

int putc (c, stream)
int c;
FILE *stream;

int putchar (c)
int c;

int fputc (c, stream)
int c;
FILE *stream;

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

### DESCRIPTION

*Putc* writes the character *c* onto the output *stream* (at the position where the file pointer, if defined, is pointing). *Putchar(c)* is defined as *putc(c, stdout)*. *Putc* and *putchar* are macros.

*Fputc* behaves like *putc*, but is a function rather than a macro. *Fputc* runs more slowly than *putc*, but it takes less space per invocation and its name can be passed as an argument to a function.

*Putw* writes the word (i.e., integer) *w* to the output *stream* (at the position at which the file pointer, if defined, is pointing). The size of a word is the size of an integer and varies from machine to machine. *Putw* neither assumes nor causes special alignment in the file.

Output streams, with the exception of the standard error stream *stderr*, are by default buffered if the output refers to a file and line-buffered if the output refers to a terminal. The standard error output stream *stderr* is by default unbuffered, but use of *freopen* [see *fopen(3S)*] will cause it to become buffered or line-buffered. When an output stream is unbuffered, information is queued for writing on the destination file or terminal as soon as written; when it is buffered, many characters are saved up and written as a block. When it is line-buffered, each line of output is queued for writing on the destination terminal as soon as the line is completed (that is, as soon as a new-line character is written or terminal input is requested). *Setbuf(3S)* or *Setbuf(3S)* may be used to change the stream's buffering strategy.

### SEE ALSO

*fclose(3S)*, *ferror(3S)*, *fopen(3S)*, *fread(3S)*, *printf(3S)*, *puts(3S)*, *setbuf(3S)*.

### DIAGNOSTICS

On success, *putc*, *fputc*, and *putchar* each return the value they have written. On failure, they return the constant EOF. This will occur if the file *stream* is not open for writing or if the output file cannot be grown. *Putw* returns nonzero when an error has occurred, otherwise zero.

### BUGS

Because it is implemented as a macro, *putc* treats incorrectly a *stream* argument with side effects. In particular, **putc(c, \*f++)**; doesn't work sensibly. *Fputc* should be used instead.

Because of possible differences in word length and byte ordering, files written using *putw* are machine-dependent, and may not be read using *getw* on a different processor.



**NAME**

putenv — change or add value to environment

**SYNOPSIS**

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

**DESCRIPTION**

*String* points to a string of the form "*name=value*." *Putenv* makes the value of the environment variable *name* equal to *value* by altering an existing variable or creating a new one. In either case, the string pointed to by *string* becomes part of the environment, so altering the string will change the environment. The space used by *string* is no longer used once a new string-defining *name* is passed to *putenv*.

**DIAGNOSTICS**

*Putenv* returns non-zero if it was unable to obtain enough space via *malloc* for an expanded environment, otherwise zero.

**SEE ALSO**

exec(2), getenv(3C), malloc(3C), environ(5).

**WARNINGS**

*Putenv* manipulates the environment pointed to by *environ*, and can be used in conjunction with *getenv*. However, *envp* (the third argument to *main*) is not changed.

This routine uses *malloc(3C)* to enlarge the environment.

After *putenv* is called, environmental variables are not in alphabetical order.

A potential error is to call *putenv* with an automatic variable as the argument, then exit the calling function while *string* is still part of the environment.

## PUTPWENT(3C)

### NAME

putpwent — write password file entry

### SYNOPSIS

```
#include <pwd.h>
int putpwent (p, f)
struct passwd *p;
FILE *f;
```

### DESCRIPTION

*Putpwent* is the inverse of *getpwent*(3C). Given a pointer to a *passwd* structure created by *getpwent* (or *getpwuid* or *getpwnam*), *putpwent* writes a line on the stream *f*, which matches the format of */etc/passwd*.

### DIAGNOSTICS

*Putpwent* returns non-zero if an error was detected during its operation, otherwise zero.

### SEE ALSO

*getpwent*(3C).

### WARNING

The above routine uses *<stdio.h>*, which causes it to increase the size of programs, not otherwise using standard I/O, more than might be expected.

**NAME**

`puts`, `fputs` -- put a string on a stream

**SYNOPSIS**

```
#include <stdio.h>

int puts (s)
char *s;

int fputs (s, stream)
char *s;
FILE *stream;
```

**DESCRIPTION**

*Puts* writes the null-terminated string pointed to by *s*, followed by a new-line character, to the standard output stream *stdout*.

*Fputs* writes the null-terminated string pointed to by *s* to the named output *stream*.

Neither function writes the terminating null character.

**DIAGNOSTICS**

Both routines return EOF on error. This will happen if the routines try to write on a file that has not been opened for writing.

**SEE ALSO**

`ferror(3S)`, `fopen(3S)`, `fread(3S)`, `printf(3S)`, `putc(3S)`.

**NOTES**

*Puts* appends a new-line character while *fputs* does not.

## QSORT(3C)

### NAME

qsort — quicker sort

### SYNOPSIS

```
void qsort ((char *) base, nel, sizeof (*base), compar)
unsigned nel;
int (*compar) ( );
```

### DESCRIPTION

*Qsort* is an implementation of the quicker-sort algorithm. It sorts a table of data in place.

*Base* points to the element at the base of the table. *Nel* is the number of elements in the table. *Compar* is the name of the comparison function, which is called with two arguments that point to the elements being compared. As the function must return an integer less than, equal to, or greater than zero, so must the first argument to be considered be less than, equal to, or greater than the second.

### NOTES

The pointer to the base of the table should be of type pointer-to-element, and cast to type pointer-to-character.

The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

The order in the output of two items which compare as equal is unpredictable.

### SEE ALSO

bsearch(3C), lsearch(3C), string(3C).  
sort(1) in the *UNIX System V User Reference Manual*.

### WARNING

The total size of the table (*nel* x *sizeof(\*base)*) must be less than 65536 in small and large model programs.

## NAME

rand, srand — simple random-number generator

## SYNOPSIS

```
int rand ( )  
void srand (seed)  
unsigned seed;
```

## DESCRIPTION

*Rand* uses a multiplicative congruential random-number generator with period  $2^{32}$  that returns successive pseudo-random numbers in the range from 0 to  $2^{15}-1$ .

*Srand* can be called at any time to reset the random-number generator to a random starting point. The generator is initially seeded with a value of 1.

## NOTE

The spectral properties of *rand* leave a great deal to be desired. *Drand48(3C)* provides a much better, though more elaborate, random-number generator.

## SEE ALSO

drand48(3C).

# SCANF(3S)

## NAME

scanf, fscanf, sscanf — convert formatted input

## SYNOPSIS

```
#include <stdio.h>

int scanf (format [ , pointer ] ... )
char *format;

int fscanf (stream, format [ , pointer ] ... )
FILE *stream;
char *format;

int sscanf (s, format [ , pointer ] ... )
char *s, *format;
```

## DESCRIPTION

*Scanf* reads from the standard input stream *stdin*. *Fscanf* reads from the named input *stream*. *Sscanf* reads from the character string *s*. Each function reads characters, interprets them according to a format, and stores the results in its arguments. Each expects, as arguments, a control string *format* described below, and a set of *pointer* arguments indicating where the converted input should be stored.

The control string usually contains conversion specifications, which are used to direct interpretation of input sequences. The control string may contain:

1. White-space characters (blanks, tabs, new-lines, or form-feeds) which, except in two cases described below, cause input to be read up to the next non-white-space character.
2. An ordinary character (not %), which must match the next character of the input stream.
3. Conversion specifications, consisting of the character %, an optional assignment suppressing character \*, an optional numerical maximum field width, an optional l (ell) or h indicating the size of the receiving variable, and a conversion code.

A conversion specification directs the conversion of the next input field; the result is placed in the variable pointed to by the corresponding argument, unless assignment suppression was indicated by \*. The suppression of assignment provides a way of describing an input field which is to be skipped. An input field is defined as a string of non-space characters; it extends to the next inappropriate character or until the field width, if specified, is exhausted. For all descriptors except "l" and "c", white space leading an input field is ignored.

The conversion code indicates the interpretation of the input field; the corresponding pointer argument must usually be of a restricted type. For a suppressed field, no pointer argument is given. The following conversion codes are legal:

- % a single % is expected in the input at this point; no assignment is done.
- d a decimal integer is expected; the corresponding argument should be an integer pointer.
- u an unsigned decimal integer is expected; the corresponding argument should be an unsigned integer pointer.
- o an octal integer is expected; the corresponding argument should be an integer pointer.
- x a hexadecimal integer is expected; the corresponding argument should be an integer pointer.
- e,f,g a floating point number is expected; the next field is converted accordingly and stored through the corresponding argument, which should be a pointer to a *float*. The input format for floating point numbers is an optionally signed string of digits, possibly containing a decimal point,

- followed by an optional exponent field consisting of an E or an e, followed by an optional +, -, or space, followed by an integer.
- s a character string is expected; the corresponding argument should be a character pointer pointing to an array of characters large enough to accept the string and a terminating \0, which will be added automatically. The input field is terminated by a white-space character.
  - c a character is expected; the corresponding argument should be a character pointer. The normal skip over white space is suppressed in this case; to read the next non-space character, use %1s. If a field width is given, the corresponding argument should refer to a character array; the indicated number of characters is read.
  - [ indicates string data and the normal skip over leading white space is suppressed. The left bracket is followed by a set of characters, which we will call the *scanset*, and a right bracket; the input field is the maximal sequence of input characters consisting entirely of characters in the scanset. The circumflex (^), when it appears as the first character in the scanset, serves as a complement operator and redefines the scanset as the set of all characters *not* contained in the remainder of the scanset string. There are some conventions used in the construction of the scanset. A range of characters may be represented by the construct *first-last*, thus [0123456789] may be expressed [0-9]. Using this convention, *first* must be lexically less than or equal to *last*, or else the dash will stand for itself. The dash will also stand for itself whenever it is the first or the last character in the scanset. To include the right square bracket as an element of the scanset, it must appear as the first character (possibly preceded by a circumflex) of the scanset, and in this case it will not be syntactically interpreted as the closing bracket. The corresponding argument must point to a character array large enough to hold the data field and the terminating \0, which will be added automatically. At least one character must match for this conversion to be considered successful.

The conversion characters **d**, **u**, **o**, and **x** may be preceded by **l** or **b** to indicate that a pointer to **long** or to **short** rather than to **int** is in the argument list. Similarly, the conversion characters **e**, **f**, and **g** may be preceded by **l** to indicate that a pointer to **double** rather than to **float** is in the argument list. The **l** or **b** modifier is ignored for other conversion characters.

*Scanf* conversion terminates at **EOF**, at the end of the control string, or when an input character conflicts with the control string. In the latter case, the offending character is left unread in the input stream.

*Scanf* returns the number of successfully matched and assigned input items; this number can be zero in the event of an early conflict between an input character and the control string. If the input ends before the first conflict or conversion, **EOF** is returned.

#### EXAMPLES

The call:

```
int i, n; float x; char name[50];
n = scanf("%d%f%s", &i, &x, name);
```

with the input line:

```
25 54.32E-1 thompson
```

will assign to *n* the value 3, to *i* the value 25, to *x* the value 5.432, and *name* will contain **thompson\0**. Or:

```
int i; float x; char name[50];
(void) scanf("%2d%f%*d%[0-9]", &i, &x, name);
```

## SCANF(3S)

with input:

56789 0123 56a72

will assign **56** to *i*, **789.0** to *x*, skip **0123**, and place the string **56\0** in *name*.  
The next call to *getchar* [see *getc(3S)*] will return **a**.

### SEE ALSO

*getc(3S)*, *printf(3S)*, *strtod(3C)*, *strtol(3C)*.

### NOTE

Trailing white space (including a new-line) is left unread unless matched in the control string.

### DIAGNOSTICS

These functions return **EOF** on end of input and a short count for missing or illegal data items.

### BUGS

The success of literal matches and suppressed assignments is not directly determinable.



## NAME

setbuf, setvbuf — assign buffering to a stream

## SYNOPSIS

```
#include <stdio.h>

void setbuf (stream, buf)
FILE *stream;
char *buf;

int setvbuf (stream, buf, type, size)
FILE *stream;
char *buf;
int type, size;
```

## DESCRIPTION

*Setbuf* may be used after a stream has been opened but before it is read or written. It causes the array pointed to by *buf* to be used instead of an automatically allocated buffer. If *buf* is the NULL pointer input/output will be completely unbuffered.

A constant **BUFSIZ**, defined in the `<stdio.h>` header file, tells how big an array is needed:

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

*Setvbuf* may be used after a stream has been opened but before it is read or written. *Type* determines how *stream* will be buffered. Legal values for *type* (defined in `stdio.h`) are:

<code>_IOFBF</code>	causes input/output to be fully buffered.
<code>_IOLBF</code>	causes output to be line buffered; the buffer will be flushed when a new-line is written, the buffer is full, or input is requested.
<code>_IONBF</code>	causes input/output to be completely unbuffered.

If *buf* is not the NULL pointer, the array it points to will be used for buffering, instead of an automatically allocated buffer. *Size* specifies the size of the buffer to be used. The constant **BUFSIZ** in `<stdio.h>` is suggested as a good buffer size. If input/output is unbuffered, *buf* and *size* are ignored.

By default, output to a terminal is line-buffered and all other input/output is fully buffered.

## SEE ALSO

`fopen(3S)`, `getc(3S)`, `malloc(3C)`, `putc(3S)`, `stdio(3S)`.

## DIAGNOSTICS

If an illegal value for *type* or *size* is provided, *setvbuf* returns a non-zero value. Otherwise, the value returned will be zero.

## NOTE

A common source of error is allocating buffer space as an “automatic” variable in a code block, and then failing to close the stream in the same block.

## SETJMP(3C)

### NAME

setjmp, longjmp — non-local goto

### SYNOPSIS

```
#include <setjmp.h>
int setjmp (env)
jmp_buf env;
void longjmp (env, val)
jmp_buf env;
int val;
```

### DESCRIPTION

These functions are useful for dealing with errors and interrupts encountered in a low-level subroutine of a program.

*Setjmp* saves its stack environment in *env* (whose type, *jmp\_buf*, is defined in the *<setjmp.h>* header file) for later use by *longjmp*. It returns the value 0.

*Longjmp* restores the environment saved by the last call of *setjmp* with the corresponding *env* argument. After *longjmp* is completed, program execution continues as if the corresponding call of *setjmp* (which must not itself have returned in the interim) had just returned the value *val*. *Longjmp* cannot cause *setjmp* to return the value 0. If *longjmp* is invoked with a second argument of 0, *setjmp* will return 1. All accessible data had values as of the time *longjmp* was called.

### SEE ALSO

signal(2).

### WARNING

If *longjmp* is called even though *env* was never primed by a call to *setjmp*, or when the last such call was in a function which has since returned, absolute chaos is guaranteed.

**NAME**

sleep — suspend execution for interval

**SYNOPSIS**

**unsigned sleep (seconds)**  
**unsigned seconds;**

**DESCRIPTION**

The current process is suspended from execution for the number of *seconds* specified by the argument. The actual suspension time may be less than that requested for two reasons: (1) Because scheduled wakeups occur at fixed 1-second intervals, (on the second, according to an internal clock) and (2) because any caught signal will terminate the *sleep* following execution of that signal's catching routine. Also, the suspension time may be longer than requested by an arbitrary amount due to the scheduling of other activity in the system. The value returned by *sleep* will be the "unslept" amount (the requested time minus the time actually slept) in case the caller had an alarm set to go off earlier than the end of the requested *sleep* time, or premature arousal due to another caught signal.

The routine is implemented by setting an alarm signal and pausing until it (or some other signal) occurs. The previous state of the alarm signal is saved and restored. The calling program may have set up an alarm signal before calling *sleep*. If the *sleep* time exceeds the time till such alarm signal, the process sleeps only until the alarm signal would have occurred. The caller's alarm catch routine is executed just before the *sleep* routine returns. But if the *sleep* time is less than the time till such alarm, the prior alarm time is reset to go off at the same time it would have without the intervening *sleep*.

**SEE ALSO**

alarm(2), pause(2), signal(2).

## SSIGNAL(3C)

### NAME

*ssignal*, *gsignal* – software signals

### SYNOPSIS

```
#include <signal.h>

int (*ssignal (sig, action))( )
int sig, (*action)( );

int gsignal (sig)
int sig;
```

### DESCRIPTION

*Ssignal* and *gsignal* implement a software facility similar to *signal(2)*. This facility is used by the Standard C Library to enable users to indicate the disposition of error conditions, and is also made available to users for their own purposes.

Software signals made available to users are associated with integers in the inclusive range 1 through 15. A call to *ssignal* associates a procedure, *action*, with the software signal *sig*; the software signal, *sig*, is raised by a call to *gsignal*. Raising a software signal causes the action established for that signal to be taken.

The first argument to *ssignal* is a number identifying the type of signal for which an action is to be established. The second argument defines the action; it is either the name of a (user-defined) *action function* or one of the manifest constants *SIG\_DFL* (default) or *SIG\_IGN* (ignore). *Ssignal* returns the action previously established for that signal type; if no action has been established or the signal number is illegal, *ssignal* returns *SIG\_DFL*.

*Gsignal* raises the signal identified by its argument, *sig*:

If an action function has been established for *sig*, then that action is reset to *SIG\_DFL* and the action function is entered with argument *sig*. *Gsignal* returns the value returned to it by the action function.

If the action for *sig* is *SIG\_IGN*, *gsignal* returns the value 1 and takes no other action.

If the action for *sig* is *SIG\_DFL*, *gsignal* returns the value 0 and takes no other action.

If *sig* has an illegal value or no action was ever specified for *sig*, *gsignal* returns the value 0 and takes no other action.

### SEE ALSO

*signal(2)*.

### NOTES

There are some additional signals with numbers outside the range 1 through 15 which are used by the Standard C Library to indicate error conditions. Thus, some signal numbers outside the range 1 through 15 are legal, although their use may interfere with the operation of the Standard C Library.

**NAME**

stdio — standard buffered input/output package

**SYNOPSIS**

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

```
FILE *stdin, *stdout, *stderr;
```

**DESCRIPTION**

The functions described in the entries of sub-class 3S of this manual constitute an efficient, user-level I/O buffering scheme. The in-line macros *getc*(3S) and *putc*(3S) handle characters quickly. The macros *getchar* and *putchar*, and the higher-level routines *fgetc*, *fgets*, *sprintf*, *fputc*, *fputs*, *fread*, *fscanf*, *fwrite*, *gets*, *getw*, *printf*, *puts*, *putw*, and *scanf* all use or act as if they use *getc* and *putc*; they can be freely intermixed.

A file with associated buffering is called a *stream* and is declared to be a pointer to a defined type **FILE**. *Fopen*(3S) creates certain descriptive data for a stream and returns a pointer to designate the stream in all further transactions. Normally, there are three open streams with constant pointers declared in the `<stdio.h>` header file and associated with the standard open files:

```

stdin      standard input file
stdout    standard output file
stderr    standard error file

```

A constant **NULL** (0) designates a nonexistent pointer.

An integer-constant **EOF** (-1) is returned upon end-of-file or error by most integer functions that deal with streams (see the individual descriptions for details).

An integer constant **BUFSIZ** specifies the size of the buffers used by the particular implementation.

Any program that uses this package must include the header file of pertinent macro definitions, as follows:

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

The functions and constants mentioned in the entries of sub-class 3S of this manual are declared in that header file and need no further declaration. The constants and the following "functions" are implemented as macros (redeclaration of these names is perilous): *getc*, *getchar*, *putc*, *putchar*, *feof*, *clearerr*, and *fileno*.

**SEE ALSO**

*open*(2), *close*(2), *lseek*(2), *pipe*(2), *read*(2), *write*(2), *ctermid*(3S), *cuserid*(3S), *fclose*(3S), *error*(3S), *fopen*(3S), *fread*(3S), *fseek*(3S), *getc*(3S), *gets*(3S), *popen*(3S), *printf*(3S), *putc*(3S), *puts*(3S), *scanf*(3S), *setbuf*(3S), *system*(3S), *tmpfile*(3S), *tmpnam*(3S), *ungetc*(3S).

**DIAGNOSTICS**

Invalid *stream* pointers will usually cause grave disorder, possibly including program termination. Individual function descriptions describe the possible error conditions.

## STDIPC(3C)

### NAME

*ftok* — standard interprocess communication package

### SYNOPSIS

```
#include <sys/types.h>
```

```
#include <sys/ipc.h>
```

```
key_t ftok(path, id)
```

```
char *path;
```

```
char id;
```

### DESCRIPTION

All interprocess communication facilities require the user to supply a key to be used by the *msgget(2)*, *semget(2)*, and *shmget(2)* system calls to obtain interprocess communication identifiers. One suggested method for forming a key is to use the *ftok* subroutine described below. Another way to compose keys is to include the project ID in the most significant byte and to use the remaining portion as a sequence number. There are many other ways to form keys, but it is necessary for each system to define standards for forming them. If some standard is not adhered to, it will be possible for unrelated processes to unintentionally interfere with each other's operation. Therefore, it is strongly suggested that the most significant byte of a key in some sense refer to a project so that keys do not conflict across a given system.

*Ftok* returns a key based on *path* and *id* that is usable in subsequent *msgget*, *semget*, and *shmget* system calls. *Path* must be the path name of an existing file that is accessible to the process. *Id* is a character which uniquely identifies a project. Note that *ftok* will return the same key for linked files when called with the same *id* and that it will return different keys when called with the same file name but different *ids*.

### SEE ALSO

*intro(2)*, *msgget(2)*, *semget(2)*, *shmget(2)*.

### DIAGNOSTICS

*Ftok* returns (**key\_t**) **-1** if *path* does not exist or if it is not accessible to the process.

### WARNING

If the file whose *path* is passed to *ftok* is removed when keys still refer to the file, future calls to *ftok* with the same *path* and *id* will return an error. If the same file is recreated, then *ftok* is likely to return a different key than it did the original time it was called.

## NAME

*strcat*, *strncat*, *strcmp*, *strncmp*, *strcpy*, *strncpy*, *strlen*, *strchr*, *strrchr*, *strpbrk*, *strspn*, *strcspn*, *strtok* — string operations

## SYNOPSIS

```
#include <string.h>

char *strcat (s1, s2)
char *s1, *s2;

char *strncat (s1, s2, n)
char *s1, *s2;
int n;

int strcmp (s1, s2)
char *s1, *s2;

int strncmp (s1, s2, n)
char *s1, *s2;
int n;

char *strcpy (s1, s2)
char *s1, *s2;

char *strncpy (s1, s2, n)
char *s1, *s2;
int n;

int strlen (s)
char *s;

char *strcbr (s, c)
char *s;
int c;

char *strrchr (s, c)
char *s;
int c;

char *strpbrk (s1, s2)
char *s1, *s2;

int strspn (s1, s2)
char *s1, *s2;

int strcspn (s1, s2)
char *s1, *s2;

char *strtok (s1, s2)
char *s1, *s2;
```

## DESCRIPTION

The arguments *s1*, *s2* and *s* point to strings (arrays of characters terminated by a null character). The functions *strcat*, *strncat*, *strcpy*, and *strncpy* all alter *s1*. These functions do not check for overflow of the array pointed to by *s1*.

*Strcat* appends a copy of string *s2* to the end of string *s1*. *Strncat* appends at most *n* characters. Each returns a pointer to the null-terminated result.

*Strcmp* compares its arguments and returns an integer less than, equal to, or greater than 0, according as *s1* is lexicographically less than, equal to, or greater than *s2*. *Strncmp* makes the same comparison but looks at most *n* characters.

*Strcpy* copies string *s2* to *s1*, stopping after the null character has been copied. *Strncpy* copies exactly *n* characters, truncating *s2* or adding null characters to *s1* if necessary. The result will not be null-terminated if the length of *s2* is *n*

## STRING(3C)

or more. Each function returns *s1*.

*Strlen* returns the number of characters in *s*, not including the terminating null character.

*Strchr* (*strrchr*) returns a pointer to the first (last) occurrence of character *c* in string *s*, or a NULL pointer if *c* does not occur in the string. The null character terminating a string is considered to be part of the string.

*Strpbrk* returns a pointer to the first occurrence in string *s1* of any character from string *s2*, or a NULL pointer if no character from *s2* exists in *s1*.

*Strspn* (*strcspn*) returns the length of the initial segment of string *s1* which consists entirely of characters from (not from) string *s2*.

*Strtok* considers the string *s1* to consist of a sequence of zero or more text tokens separated by spans of one or more characters from the separator string *s2*. The first call (with pointer *s1* specified) returns a pointer to the first character of the first token, and will have written a null character into *s1* immediately following the returned token. The function keeps track of its position in the string between separate calls, so that subsequent calls (which must be made with the first argument a NULL pointer) will work through the string *s1* immediately following that token. In this way subsequent calls will work through the string *s1* until no tokens remain. The separator string *s2* may be different from call to call. When no token remains in *s1*, a NULL pointer is returned.

### NOTE

For user convenience, all these functions are declared in the optional *<string.h>* header file.

### BUGS

*Strcmp* and *strncmp* use native character comparison, which is signed on PDP-11s and VAX-11s, unsigned on other machines. Thus the sign of the value returned when one of the characters has its high-order bit set is implementation-dependent.

Character movement is performed differently in different implementations. Thus overlapping moves may yield surprises.



## NAME

strtod, atof — convert string to double-precision number

## SYNOPSIS

```
double strtod (str, ptr)
```

```
char *str, **ptr;
```

```
double atof (str)
```

```
char *str;
```

## DESCRIPTION

*Strtod* returns as a double-precision floating-point number the value represented by the character string pointed to by *str*. The string is scanned up to the first unrecognized character.

*Strtod* recognizes an optional string of "white-space" characters [as defined by *isspace* in *ctype*(3C)], then an optional sign, then a string of digits optionally containing a decimal point, then an optional *e* or *E* followed by an optional sign or space, followed by an integer.

If the value of *ptr* is not (char \*\*)NULL, a pointer to the character terminating the scan is returned in the location pointed to by *ptr*. If no number can be formed, *\*ptr* is set to *str*, and zero is returned.

*Atof(str)* is equivalent to *strtod(str, (char \*\*)NULL)*.

## SEE ALSO

*ctype*(3C), *scanf*(3S), *strtoul*(3C).

## DIAGNOSTICS

If the correct value would cause overflow,  $\pm$ HUGE is returned (according to the sign of the value), and *errno* is set to ERANGE.

If the correct value would cause underflow, zero is returned and *errno* is set to ERANGE.

53

## STRTOL(3C)

### NAME

`strtol`, `atol`, `atoi` — convert string to integer

### SYNOPSIS

```
long strtol (str, ptr, base)
char *str, **ptr;
int base;

long atol (str)
char *str;

int atoi (str)
char *str;
```

### DESCRIPTION

*Strtol* returns as a long integer the value represented by the character string pointed to by *str*. The string is scanned up to the first character inconsistent with the base. Leading "white-space" characters [as defined by *isspace* in *ctype*(3C)] are ignored.

If the value of *ptr* is not (char \*\*)NULL, a pointer to the character terminating the scan is returned in the location pointed to by *ptr*. If no integer can be formed, that location is set to *str*, and zero is returned.

If *base* is positive (and not greater than 36), it is used as the base for conversion. After an optional leading sign, leading zeros are ignored, and "0x" or "0X" is ignored if *base* is 16.

If *base* is zero, the string itself determines the base thusly: After an optional leading sign a leading zero indicates octal conversion, and a leading "0x" or "0X" hexadecimal conversion. Otherwise, decimal conversion is used.

Truncation from long to int can, of course, take place upon assignment or by an explicit cast.

*Atol*(*str*) is equivalent to *strtol*(*str*, (char \*\*)NULL, 10).

*Atoi*(*str*) is equivalent to (int) *strtol*(*str*, (char \*\*)NULL, 10).

### SEE ALSO

`ctype`(3C), `scanf`(3S), `strtod`(3C).

### BUGS

Overflow conditions are ignored.

**NAME**

swab — swap bytes

**SYNOPSIS**

```
void swab (from, to, nbytes)
char *from, *to;
int nbytes;
```

**DESCRIPTION**

*Swab* copies *nbytes* bytes pointed to by *from* to the array pointed to by *to*, exchanging adjacent even and odd bytes. It is useful for carrying binary data between PDP-11s and other machines. *Nbytes* should be even and non-negative. If *nbytes* is odd and positive *swab* uses *nbytes*-1 instead. If *nbytes* is negative, *swab* does nothing.

## SYSTEM(3S)

### NAME

system — issue a shell command

### SYNOPSIS

```
#include <stdio.h>
int system (string)
char *string;
```

### DESCRIPTION

*System* causes the *string* to be given to *sh*(1) as input, as if the string had been typed as a command at a terminal. The current process waits until the shell has completed, then returns the exit status of the shell.

### FILES

*/bin/sh*

### SEE ALSO

*exec*(2),  
*sh*(1) in the Runtime System manual.

### DIAGNOSTICS

*System* forks to create a child process that in turn *exec*'s */bin/sh* in order to execute *string*. If the fork or *exec* fails, *system* returns a negative value and sets *errno*.

**NAME**  
tmpfile – create a temporary file

**SYNOPSIS**  
`#include <stdio.h>`  
`FILE *tmpfile ()`

**DESCRIPTION**  
*Tmpfile* creates a temporary file using a name generated by *tmpnam*(3S), and returns a corresponding FILE pointer. If the file cannot be opened, an error message is printed using *perror*(3C), and a NULL pointer is returned. The file will automatically be deleted when the process using it terminates. The file is opened for update ("w+").

**SEE ALSO**  
*creat*(2), *unlink*(2), *fopen*(3S), *mktemp*(3C), *perror*(3C), *tmpnam*(3S).

# TMPNAM(3S)

## NAME

`tmpnam`, `tmpnam` – create a name for a temporary file

## SYNOPSIS

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

```
char *tmpnam (s)
```

```
char *s;
```

```
char *tmpnam (dir, pfx)
```

```
char *dir, *pfx;
```

## DESCRIPTION

These functions generate file names that can safely be used for a temporary file.

*Tmpnam* always generates a file name using the path-prefix defined as `P_tmpdir` in the `<stdio.h>` header file. If *s* is NULL, *tmpnam* leaves its result in an internal static area and returns a pointer to that area. The next call to *tmpnam* will destroy the contents of the area. If *s* is not NULL, it is assumed to be the address of an array of at least `L_tmpnam` bytes, where `L_tmpnam` is a constant defined in `<stdio.h>`; *tmpnam* places its result in that array and returns *s*.

*Tempnam* allows the user to control the choice of a directory. The argument *dir* points to the name of the directory in which the file is to be created. If *dir* is NULL or points to a string which is not a name for an appropriate directory, the path-prefix defined as `P_tmpdir` in the `<stdio.h>` header file is used. If that directory is not accessible, `/tmp` will be used as a last resort. This entire sequence can be up-staged by providing an environment variable `TMPDIR` in the user's environment, whose value is the name of the desired temporary-file directory.

Many applications prefer their temporary files to have certain favorite initial letter sequences in their names. Use the *pfx* argument for this. This argument may be NULL or point to a string of up to five characters to be used as the first few characters of the temporary-file name.

*Tempnam* uses `malloc(3C)` to get space for the constructed file name, and returns a pointer to this area. Thus, any pointer value returned from *tempnam* may serve as an argument to `free` [see `malloc(3C)`]. If *tempnam* cannot return the expected result for any reason, i.e., `malloc(3C)` failed, or none of the above mentioned attempts to find an appropriate directory was successful, a NULL pointer will be returned.

## NOTES

These functions generate a different file name each time they are called.

Files created using these functions and either `fopen(3S)` or `creat(2)` are temporary only in the sense that they reside in a directory intended for temporary use, and their names are unique. It is the user's responsibility to use `unlink(2)` to remove the file when its use is ended.

## SEE ALSO

`creat(2)`, `unlink(2)`, `fopen(3S)`, `malloc(3C)`, `mktemp(3C)`, `tmpfile(3S)`.

## BUGS

If called more than 17,576 times in a single process, these functions will start recycling previously used names.

Between the time a file name is created and the file is opened, it is possible for some other process to create a file with the same name. This can never happen if that other process is using these functions or `mktemp`, and the file names are chosen so as to render duplication by other means unlikely.

## NAME

`tsearch`, `tfind`, `tdelete`, `twalk` — manage binary search trees

## SYNOPSIS

```
#include <search.h>

char *tsearch ((char *) key, (char **) rootp, compar)
int (*compar) ( );

char *tfind ((char *) key, (char **) rootp, compar)
int (*compar) ( );

char *tdelete ((char *) key, (char **) rootp, compar)
int (*compar) ( );

void twalk ((char *) root, action)
void (*action) ( );
```

## DESCRIPTION

*Tsearch*, *tfind*, *tdelete*, and *twalk* are routines for manipulating binary search trees. They are generalized from Knuth (6.2.2) Algorithms T and D. All comparisons are done with a user-supplied routine. This routine is called with two arguments, the pointers to the elements being compared. It returns an integer less than, equal to, or greater than 0, according to whether the first argument is to be considered less than, equal to or greater than the second argument. The comparison function need not compare every byte, so arbitrary data may be contained in the elements in addition to the values being compared.

*Tsearch* is used to build and access the tree. *Key* is a pointer to a datum to be accessed or stored. If there is a datum in the tree equal to *\*key* (the value pointed to by *key*), a pointer to this found datum is returned. Otherwise, *\*key* is inserted, and a pointer to it returned. Only pointers are copied, so the calling routine must store the data. *Rootp* points to a variable that points to the root of the tree. A NULL value for the variable pointed to by *rootp* denotes an empty tree; in this case, the variable will be set to point to the datum which will be at the root of the new tree.

Like *tsearch*, *tfind* will search for a datum in the tree, returning a pointer to it if found. However, if it is not found, *tfind* will return a NULL pointer. The arguments for *tfind* are the same as for *tsearch*.

*Tdelete* deletes a node from a binary search tree. The arguments are the same as for *tsearch*. The variable pointed to by *rootp* will be changed if the deleted node was the root of the tree. *Tdelete* returns a pointer to the parent of the deleted node, or a NULL pointer if the node is not found.

*Twalk* traverses a binary search tree. *Root* is the root of the tree to be traversed. (Any node in a tree may be used as the root for a walk below that node.) *Action* is the name of a routine to be invoked at each node. This routine is, in turn, called with three arguments. The first argument is the address of the node being visited. The second argument is a value from an enumeration data type `typedef enum { preorder, postorder, endorder, leaf } VISIT`; (defined in the `<search.h>` header file), depending on whether this is the first, second or third time that the node has been visited (during a depth-first, left-to-right traversal of the tree), or whether the node is a leaf. The third argument is the level of the node in the tree, with the root being level zero.

The pointers to the key and the root of the tree should be of type pointer-to-element, and cast to type pointer-to-character. Similarly, although declared as type pointer-to-character, the value returned should be cast into type pointer-to-element.

## TSEARCH(3C)

### EXAMPLE

The following code reads in strings and stores structures containing a pointer to each string and a count of its length. It then walks the tree, printing out the stored strings and their lengths in alphabetical order.

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

struct node {          /* pointers to these are stored in the tree */
    char *string;
    int length;
};
char string_space[10000]; /* space to store strings */
struct node nodes[500]; /* nodes to store */
struct node *root = NULL; /* this points to the root */

main( )
{
    char *strptr = string_space;
    struct node *nodeptr = nodes;
    void print_node( ), twalk( );
    int i = 0, node_compare( );

    while (gets(strptr) != NULL && i++ < 500) {
        /* set node */
        nodeptr->string = strptr;
        nodeptr->length = strlen(strptr);
        /* put node into the tree */
        (void) tsearch((char *)nodeptr, &root,
            node_compare);
        /* adjust pointers, so we don't overwrite tree */
        strptr += nodeptr->length + 1;
        nodeptr++;
    }
    twalk(root, print_node);

    /*
       This routine compares two nodes, based on an
       alphabetical ordering of the string field.
    */
    int
    node_compare(node1, node2)
    struct node *node1, *node2;
    {
        return strcmp(node1->string, node2->string);
    }
    /*
       This routine prints out a node, the first time
       twalk encounters it.
    */
}
```



```

void
print_node(node, order, level)
struct node **node;
VISIT order;
int level;
{
    if (order == preorder || order == leaf) {
        (void)printf("string = %20s, length = %d\n",
            (*node)->string, (*node)->length);
    }
}

```

**SEE ALSO**

*bsearch*(3C), *hsearch*(3C), *lsearch*(3C).

**DIAGNOSTICS**

A NULL pointer is returned by *tsearch* if there is not enough space available to create a new node.

A NULL pointer is returned by *tsearch*, *tfind* and *tdelete* if **rootp** is NULL on entry.

If the datum is found, both *tsearch* and *tfind* return a pointer to it. If not, *tfind* returns NULL, and *tsearch* returns a pointer to the inserted item.

**WARNINGS**

The **root** argument to *twalk* is one level of indirection less than the **rootp** arguments to *tsearch* and *tdelete*.

There are two nomenclatures used to refer to the order in which tree nodes are visited. *Tsearch* uses *preorder*, *postorder* and *endorder* to respectively refer to visiting a node before any of its children, after its left child and before its right, and after both its children. The alternate nomenclature uses *preorder*, *inorder* and *postorder* to refer to the same visits, which could result in some confusion over the meaning of *postorder*.

**BUGS**

If the calling function alters the pointer to the root, results are unpredictable.

## TTYNAME(3C)

### NAME

`ttyname`, `isatty` — find name of a terminal

### SYNOPSIS

```
char *ttyname (fildes)
```

```
int fildes;
```

```
int isatty (fildes)
```

```
int fildes;
```

### DESCRIPTION

*Ttyname* returns a pointer to a string containing the null-terminated path name of the terminal device associated with file descriptor *fildes*.

*Isatty* returns 1 if *fildes* is associated with a terminal device, 0 otherwise.

### FILES

*/dev/\**

### DIAGNOSTICS

*Ttyname* returns a NULL pointer if *fildes* does not describe a terminal device in directory */dev*.

### BUGS

The return value points to static data whose content is overwritten by each call.

**NAME**

ttyslot — find the slot in the utmp file of the current user

**SYNOPSIS**

int ttyslot ( )

**DESCRIPTION**

*Ttyslot* returns the index of the current user's entry in the */etc/utmp* file. This is accomplished by actually scanning the file */etc/inittab* for the name of the terminal associated with the standard input, the standard output, or the error output (0, 1 or 2).

**FILES**

*/etc/inittab*  
*/etc/utmp*

**SEE ALSO**

*getut(3C)*, *ttyname(3C)*.

**DIAGNOSTICS**

A value of 0 is returned if an error was encountered while searching for the terminal name or if none of the above file descriptors is associated with a terminal device.

# UNGETC(3S)

## NAME

`ungetc` — push character back into input stream

## SYNOPSIS

```
#include <stdio.h>
int ungetc (c, stream)
int c;
FILE *stream;
```

## DESCRIPTION

*Ungetc* inserts the character *c* into the buffer associated with an input *stream*. That character, *c*, will be returned by the next *getc*(3S) call on that *stream*. *Ungetc* returns *c*, and leaves the file *stream* unchanged.

One character of pushback is guaranteed, provided something has already been read from the stream and the stream is actually buffered. In the case that *stream* is *stdin*, one character may be pushed back onto the buffer without a previous read statement.

If *c* equals EOF, *ungetc* does nothing to the buffer and returns EOF.

*Fseek*(3S) erases all memory of inserted characters.

## SEE ALSO

*fseek*(3S), *getc*(3S), *setbuf*(3S).

## DIAGNOSTICS

*Ungetc* returns EOF if it cannot insert the character.

## NAME

*vprintf*, *vfprintf*, *vsprintf* — print formatted output of a *varargs* argument list

## SYNOPSIS

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

int vprintf (format, ap)
char *format;
va_list ap;

int vfprintf (stream, format, ap)
FILE *stream;
char *format;
va_list ap;

int vsprintf (s, format, ap)
char *s, *format;
va_list ap;
```

## DESCRIPTION

*vprintf*, *vfprintf*, and *vsprintf* are the same as *printf*, *fprintf*, and *sprintf* respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by *varargs*(5).

## EXAMPLE

The following demonstrates how *vfprintf* could be used to write an error routine.

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

/*
 * error should be called like
 * error(function_name, format, arg1, arg2...);
 */
/*VARARGS0*/
void
error(va_alist)
/* Note that the function_name and format arguments cannot be
 * separately declared because of the definition of varargs.
 */
va_dcl
{
    va_list args;
    char *fmt;

    va_start(args);
    /* print out name of function causing error */
    (void)fprintf(stderr, "ERROR in %s: ", va_arg(args, char *));
    fmt = va_arg(args, char *);
    /* print out remainder of message */
    (void)vfprintf(fmt, args);
    va_end(args);
    (void)abort( );
}
}
```

## SEE ALSO

*vprintf*(3X), *varargs*(5).

2

3

4

# TABLE OF CONTENTS OF MATH SUBROUTINES

## 3M. Math Subroutines and Libraries

bessel.....	Bessel functions
erf.....	error function and complementary error function
exp.....	exponential, logarithm, power, square root functions
floor.....	floor, ceiling, remainder absolute value functions
gamma.....	log gamma function
hypot.....	Euclidean distance function
matherr.....	error-handling function
sinh.....	hyperbolic functions
trig.....	trigonometric functions

1

2

3



## NAME

$j_0$ ,  $j_1$ ,  $j_n$ ,  $y_0$ ,  $y_1$ ,  $y_n$  — Bessel functions

## SYNOPSIS

```
#include <math.h>
```

```
double j0 (x)
```

```
double x;
```

```
double j1 (x)
```

```
double x;
```

```
double jn (n, x)
```

```
int n;
```

```
double x;
```

```
double y0 (x)
```

```
double x;
```

```
double y1 (x)
```

```
double x;
```

```
double yn (n, x)
```

```
int n;
```

```
double x;
```

## DESCRIPTION

$J_0$  and  $J_1$  return Bessel functions of  $x$  of the first kind of orders 0 and 1 respectively.  $J_n$  returns the Bessel function of  $x$  of the first kind of order  $n$ .

$Y_0$  and  $Y_1$  return Bessel functions of  $x$  of the second kind of orders 0 and 1 respectively.  $Y_n$  returns the Bessel function of  $x$  of the second kind of order  $n$ . The value of  $x$  must be positive.

## DIAGNOSTICS

Non-positive arguments cause  $y_0$ ,  $y_1$  and  $y_n$  to return the value **-HUGE** and to set *errno* to **EDOM**. In addition, a message indicating **DOMAIN** error is printed on the standard error output.

Arguments too large in magnitude cause  $j_0$ ,  $j_1$ ,  $y_0$  and  $y_1$  to return zero and to set *errno* to **ERANGE**. In addition, a message indicating **TLOSS** error is printed on the standard error output.

These error-handling procedures may be changed with the function *matherr*(3M).

## SEE ALSO

*matherr*(3M).

S-3M

## ERF(3M)

### NAME

erf, erfc – error function and complementary error function

### SYNOPSIS

```
#include <math.h>
double erf (x)
double x;
double erfc (x)
double x;
```

### DESCRIPTION

*Erf* returns the error function of  $x$ , defined as  $\frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$ .

*Erfc*, which returns  $1.0 - erf(x)$ , is provided because of the extreme loss of relative accuracy if *erf*( $x$ ) is called for large  $x$  and the result subtracted from 1.0 (e.g., for  $x = 5$ , 12 places are lost).

### SEE ALSO

exp(3M).

## NAME

*exp*, *log*, *log10*, *pow*, *sqrt* — exponential, logarithm, power, square root functions

## SYNOPSIS

```
#include <math.h>
```

```
double exp (x)
```

```
double x;
```

```
double log (x)
```

```
double x;
```

```
double log10 (x)
```

```
double x;
```

```
double pow (x, y)
```

```
double x, y;
```

```
double sqrt (x)
```

```
double x;
```

## DESCRIPTION

*Exp* returns  $e^x$ .

*Log* returns the natural logarithm of  $x$ . The value of  $x$  must be positive.

*Log10* returns the logarithm base ten of  $x$ . The value of  $x$  must be positive.

*Pow* returns  $x^y$ . If  $x$  is zero,  $y$  must be positive. If  $x$  is negative,  $y$  must be an integer.

*Sqrt* returns the non-negative square root of  $x$ . The value of  $x$  may not be negative.

## DIAGNOSTICS

*Exp* returns HUGE when the correct value would overflow, or 0 when the correct value would underflow, and sets *errno* to ERANGE.

*Log* and *log10* return -HUGE and set *errno* to EDOM when  $x$  is non-positive. A message indicating DOMAIN error (or SING error when  $x$  is 0) is printed on the standard error output.

*Pow* returns 0 and sets *errno* to EDOM when  $x$  is 0 and  $y$  is non-positive, or when  $x$  is negative and  $y$  is not an integer. In these cases a message indicating DOMAIN error is printed on the standard error output. When the correct value for *pow* would overflow or underflow, *pow* returns  $\pm$ HUGE or 0 respectively, and sets *errno* to ERANGE.

*Sqrt* returns 0 and sets *errno* to EDOM when  $x$  is negative. A message indicating DOMAIN error is printed on the standard error output.

These error-handling procedures may be changed with the function *matherr*(3M).

## SEE ALSO

*hypot*(3M), *matherr*(3M), *sinh*(3M).

## FLOOR(3M)

### NAME

floor, ceil, fmod, fabs — floor, ceiling, remainder, absolute value functions

### SYNOPSIS

```
#include <math.h>
double floor (x)
double x;
double ceil (x)
double x;
double fmod (x, y)
double x, y;
double fabs (x)
double x;
```

### DESCRIPTION

*Floor* returns the largest integer (as a double-precision number) not greater than  $x$ .

*Ceil* returns the smallest integer not less than  $x$ .

*Fmod* returns the floating-point remainder of the division of  $x$  by  $y$ : zero if  $y$  is zero or if  $x/y$  would overflow; otherwise the number  $f$  with the same sign as  $x$ , such that  $x = iy + f$  for some integer  $i$ , and  $|f| < |y|$ .

*Fabs* returns the absolute value of  $x$ ,  $|x|$ .

### SEE ALSO

abs(3C).

## NAME

gamma - log gamma function

## SYNOPSIS

```
#include <math.h>

double gamma (x)
double x;

extern int signgam;
```

## DESCRIPTION

*Gamma* returns  $\ln(|\Gamma(x)|)$ , where  $\Gamma(x)$  is defined as  $\int_0^{\infty} e^{-t} t^{x-1} dt$ . The sign of  $\Gamma(x)$  is returned in the external integer *signgam*. The argument *x* may not be a non-positive integer.

The following C program fragment might be used to calculate  $\Gamma$ :

```
if ((y = gamma(x)) > LN_MAXDOUBLE)
    error();
y = signgam * exp(y);
```

where LN\_MAXDOUBLE is the least value that causes *exp*(3M) to return a range error, and is defined in the *<values.h>* header file.

## DIAGNOSTICS

For non-negative integer arguments HUGE is returned, and *errno* is set to EDOM. A message indicating SING error is printed on the standard error output.

If the correct value would overflow, *gamma* returns HUGE and sets *errno* to ERANGE.

These error-handling procedures may be changed with the function *matherr*(3M).

## SEE ALSO

*exp*(3M), *matherr*(3M), *values*(5).

## HYPOT(3M)

### NAME

hypot — Euclidean distance function

### SYNOPSIS

```
#include <math.h>
```

```
double hypot (x, y)
```

```
double x, y;
```

### DESCRIPTION

*Hypot* returns

$$\sqrt{x \cdot x + y \cdot y},$$

taking precautions against unwarranted overflows.

### DIAGNOSTICS

When the correct value would overflow, *hypot* returns **HUGE** and sets *errno* to **ERANGE**.

These error-handling procedures may be changed with the function *matherr*(3M).

### SEE ALSO

*matherr*(3M).

## NAME

`matherr` — error-handling function

## SYNOPSIS

```
#include <math.h>

int matherr (x)
struct exception *x;
```

## DESCRIPTION

*Matherr* is invoked by functions in the Math Library when errors are detected. Users may define their own procedures for handling errors, by including a function named *matherr* in their programs. *Matherr* must be of the form described above. When an error occurs, a pointer to the exception structure *x* will be passed to the user-supplied *matherr* function. This structure, which is defined in the *<math.h>* header file, is as follows:

```
struct exception {
    int type;
    char *name;
    double arg1, arg2, retval;
};
```

The element *type* is an integer describing the type of error that has occurred, from the following list of constants (defined in the header file):

DOMAIN	argument domain error
SING	argument singularity
OVERFLOW	overflow range error
UNDERFLOW	underflow range error
TLOSS	total loss of significance
PLOSS	partial loss of significance

The element *name* points to a string containing the name of the function that incurred the error. The variables *arg1* and *arg2* are the arguments with which the function was invoked. *Retval* is set to the default value that will be returned by the function unless the user's *matherr* sets it to a different value.

If the user's *matherr* function returns non-zero, no error message will be printed, and *errno* will not be set.

If *matherr* is not supplied by the user, the default error-handling procedures, described with the math functions involved, will be invoked upon error. These procedures are also summarized in the table below. In every case, *errno* is set to EDOM or ERANGE and the program continues.

## EXAMPLE

```
#include <math.h>

int
matherr(x)
register struct exception *x;
{
    switch (x->type) {
    case DOMAIN:
        /* change sqrt to return sqrt(-arg1), not 0 */
        if (!strcmp(x->name, "sqrt")) {
            x->retval = sqrt(-x->arg1);
            return (0); /* print message and set errno */
        }
    }
}
```

```

case SING:
    /* all other domain or sing errors, print message and abort */
    fprintf(stderr, "domain error in %s\n", x->name);
    abort();
case PLOSS:
    /* print detailed error message */
    fprintf(stderr, "loss of significance in %s(%g) = %g\n",
            x->name, x->arg1, x->retval);
    return (1); /* take no other action */
}
return (0); /* all other errors, execute default procedure */
}

```

**DEFAULT ERROR HANDLING PROCEDURES**

type	<i>Types of Errors</i>					
	DOMAIN	SING	OVERFLOW	UNDERFLOW	TLOSS	PLOSS
<i>errno</i>	EDOM	EDOM	ERANGE	ERANGE	ERANGE	ERANGE
BESSEL:	-	-	-	-	M, 0	*
y0, y1, yn (arg ≤ 0)	M, -H	-	-	-	-	-
EXP:	-	-	H	0	-	-
LOG, LOG10: (arg < 0) (arg = 0)	M, -H -	- M, -H	- -	- -	- -	- -
POW: neg ** non-int 0 ** non-pos	- M, 0	- -	±H -	0 -	- -	- -
SQRT:	M, 0	-	-	-	-	-
GAMMA:	-	M, H	H	-	-	-
HYPOT:	-	-	H	-	-	-
SINH:	-	-	±H	-	-	-
COSH:	-	-	H	-	-	-
SIN, COS, TAN: --	-	-	-	M, 0	*	-
ASIN, ACOS, ATAN2: M, 0	-	-	-	-	-	-

ABBREVIATIONS	
*	As much as possible of the value is returned.
M	Message is printed (EDOM error).
H	HUGE is returned.
-H	-HUGE is returned.
±H	HUGE or -HUGE is returned.
0	0 is returned.



## NAME

*sinh*, *cosh*, *tanh* — hyperbolic functions

## SYNOPSIS

```
#include <math.h>
```

```
double sinh (x)
```

```
double x;
```

```
double cosh (x)
```

```
double x;
```

```
double tanh (x)
```

```
double x;
```

## DESCRIPTION

*Sinh*, *cosh*, and *tanh* return, respectively, the hyperbolic sine, cosine and tangent of their argument.

## DIAGNOSTICS

*Sinh* and *cosh* return HUGE (and *sinh* may return -HUGE for negative *x*) when the correct value would overflow and set *errno* to ERANGE.

These error-handling procedures may be changed with the function *matherr*(3M).

## SEE ALSO

*matherr*(3M).

## TRIG(3M)

### NAME

sin, cos, tan, asin, acos, atan, atan2 — trigonometric functions

### SYNOPSIS

```
#include <math.h>

double sin (x)
double x;

double cos (x)
double x;

double tan (x)
double x;

double asin (x)
double x;

double acos (x)
double x;

double atan (x)
double x;

double atan2 (y, x)
double y, x;
```

### DESCRIPTION

*Sin*, *cos* and *tan* return respectively the sine, cosine and tangent of their argument,  $x$ , measured in radians.

*Asin* returns the arcsine of  $x$ , in the range  $-\pi/2$  to  $\pi/2$ .

*Acos* returns the arccosine of  $x$ , in the range 0 to  $\pi$ .

*Atan* returns the arctangent of  $x$ , in the range  $-\pi/2$  to  $\pi/2$ .

*Atan2* returns the arctangent of  $y/x$ , in the range  $-\pi$  to  $\pi$ , using the signs of both arguments to determine the quadrant of the return value.

### DIAGNOSTICS

*Sin*, *cos*, and *tan* lose accuracy when their argument is far from zero. For arguments sufficiently large, these functions return zero when there would otherwise be a complete loss of significance. In this case a message indicating TLOSS error is printed on the standard error output. For less extreme arguments causing partial loss of significance, a PLOSS error is generated but no message is printed. In both cases, *errno* is set to **ERANGE**.

If the magnitude of the argument of *asin* or *acos* is greater than one, or if both arguments of *atan2* are zero, zero is returned and *errno* is set to **EDOM**. In addition, a message indicating DOMAIN error is printed on the standard error output.

These error-handling procedures may be changed with the function *matherr*(3M).

### SEE ALSO

*matherr*(3M).

## TABLE OF CONTENTS OF MISCELLANEOUS SUBROUTINES

### 3X. Various Specialized Libraries

assert.....	verify program assertion
cursor.....	CRT screen handling and optimization package
ldahread.....	read the archive header of a member of an archive file
ldclose.....	close a common object file
ldfread.....	read the file header of a common object file
ldgetname.....	retrieve symbol name for common object file symbol table entry
ldline.....	manipulate line number entries of a common object file function
ldlseek.....	seek to line number entries of a section of a common object file
ldahseek.....	seek to the optional file header of a common object file
ldopen.....	open a common object file for reading
ldrseek.....	seek to relocation entries of a section of a common object file
ldsread.....	read an indexed/named section header of a common object file
ldsseek.....	seek to an indexed/named section of a common object file
ldtindex.....	compute the index of a symbol table entry of a common object file
ldtread.....	read an indexed symbol table entry of a common object file
ldtseek.....	seek to the symbol table of a common object file
logname.....	return login name of user
malloc.....	fast main memory allocator
plot.....	graphics interface subroutines
regcmp.....	compile and execute regular expression
spul.....	access long integer data in a machine-independent fashion
vprintf.....	print formatted output of a varargs argument list

1

2

3

**NAME**

assert — verify program assertion

**SYNOPSIS**

```
#include <assert.h>
```

```
assert (expression)
```

```
int expression;
```

**DESCRIPTION**

This macro is useful for putting diagnostics into programs. When it is executed, if *expression* is false (zero), *assert* prints

“Assertion failed: *expression*, file *xyz*, line *nnn*”

on the standard error output and aborts. In the error message, *xyz* is the name of the source file and *nnn* the source line number of the *assert* statement.

Compiling with the preprocessor option `-DNDEBUG` [see *cpp*(1)], or with the preprocessor control statement `#define NDEBUG` ahead of the `#include <assert.h>` statement, will stop assertions from being compiled into the program.

**SEE ALSO**

*abort*(3C).

*cpp*(1) in the Runtime System manual.

# CURSES(3X)

## NAME

curses — CRT screen handling and optimization package

## SYNOPSIS

```
#include <curses.h>
cc [ flags ] files -lcurses [ libraries ]
```

## DESCRIPTION

These routines give the user a method of updating screens with reasonable optimization. In order to initialize the routines, the routine *initscr()* must be called before any of the other routines that deal with windows and screens are used. The routine *endwin()* should be called before exiting. To get character-at-a-time input without echoing, (most interactive, screen oriented-programs want this) after calling *initscr()* you should call "*nonl(); cbreak(); noecho();*"

The full curses interface permits manipulation of data structures called *windows* which can be thought of as two dimensional arrays of characters representing all or part of a CRT screen. A default window called *stdscr* is supplied, and others can be created with *newwin*. Windows are referred to by variables declared "WINDOW \*", the type WINDOW is defined in curses.h to be a C structure. These data structures are manipulated with functions described below, among which the most basic are *move*, and *addch*. (More general versions of these functions are included with names beginning with 'w', allowing you to specify a window. The routines not beginning with 'w' affect *stdscr*.) Then *refresh()* is called, telling the routines to make the users CRT screen look like *stdscr*.

Mini-Curses is a subset of curses which does not allow manipulation of more than one window. To invoke this subset, use -DMINICURSES as a cc option. This level is smaller and faster than full curses.

If the environment variable TERMINFO is defined, any program using curses will check for a local terminal definition before checking in the standard place. For example, if the standard place is */usr/lib/terminfo*, and TERM is set to "vt100", then normally the compiled file is found in */usr/lib/terminfo/v/vt100*. (The "v" is copied from the first letter of "vt100" to avoid creation of huge directories.) However, if TERMINFO is set to */usr/mark/myterms*, curses will first check */opusr/mark/myterms/v/vt100*, and if that fails, will then check */usr/lib/terminfo/v/vt100*. This is useful for developing experimental definitions or when write permission in */usr/lib/terminfo* is not available.

## SEE ALSO

*terminfo(4)*.

## FUNCTIONS

Routines listed here may be called when using the full curses. Those marked with an asterisk may be called when using Mini-Curses.

<i>addch(ch)*</i>	add a character to <i>stdscr</i> (like <i>putchar</i> ) (wraps to next line at end of line)
<i>addstr(str)*</i>	calls <i>addch</i> with each character in <i>str</i>
<i>attroff(attrs)*</i>	turn off attributes named
<i>attron(attrs)*</i>	turn on attributes named
<i>attrset(attrs)*</i>	set current attributes to <i>attrs</i>
<i>baudrate()*</i>	current terminal speed
<i>beep()*</i>	sound beep on terminal
<i>box(win, vert, hor)</i>	draw a box around edges of <i>win</i> <i>vert</i> and <i>hor</i> are chars to use for <i>vert</i> . and <i>hor</i> . edges of box
<i>clear()</i>	clear <i>stdscr</i>

<p>clearok(win, bf)  clrrobot()  clrtoeol()  cbreak()*  delay_output(ms)*  delch()  deleteln()  delwin(win)  doupdate()  echo()*  endwin()*  erase()  erasecliar()  fixterm()  flash()  flushinp()*  getch()*  getstr(str)  gettmode()  getyx(win, y, x)  has_ic()  has_il()  idlok(win, bf)*  inch()  initscr()*  insch(c)  insertln()  intrflush(win, bf)  keypad(win, bf)  killchar()  leaveok(win, flag)</p> <p>longname()  meta(win, flag)*  move(y, x)*  mvaddch(y, x, ch)  mvaddstr(y, x, str)  mvcur(olddrow, oldcol, newrow, newcol)</p> <p>mvdelch(y, x)  mvgetch(y, x)  mvgetstr(y, x)  mvinch(y, x)  mvinsch(y, x, c)  mvprintw(y, x, fmt, args)  mvscanw(y, x, fmt, args)  mwaddch(win, y, x, ch)  mwaddstr(win, y, x, str)  mwdelch(win, y, x)  mwgetch(win, y, x)  mwgetstr(win, y, x)  mwwin(win, by, bx)  mwwinch(win, y, x)  mwwinsch(win, y, x, c)  mwprintw(win, y, x, fmt, args)  mwscanw(win, y, x, fmt, args)</p>	<p>clear screen before next redraw of <i>win</i>  clear to bottom of <i>stdscr</i>  clear to end of line on <i>stdscr</i>  set cbreak mode  insert ms millisecond pause in output  delete a character  delete a line  delete <i>win</i>  update screen from all wnooutrefresh  set echo mode  end window modes  erase <i>stdscr</i>  return user's erase character  restore tty to "in curses" state  flash screen or beep  throw away any typeahead  get a char from tty  get a string through <i>stdscr</i>  establish current tty modes  get (y, x) co-ordinates  true if terminal can do insert character  true if terminal can do insert line  use terminal's insert/delete line if bf != 0  get char at current (y, x) co-ordinates  initialize screens  insert a char  insert a line  interrupts flush output if bf is TRUE  enable keypad input  return current user's kill character  OK to leave cursor anywhere after refresh if  flag!=0 for <i>win</i>, otherwise cursor must be left  at current position.  return verbose name of terminal  allow meta characters on input if flag != 0  move to (y, x) on <i>stdscr</i>  move(y, x) then addch(ch)  similar...</p> <p>low level cursor motion  like delch, but move(y, x) first  etc.</p>
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

## CURSES(3X)

`newpad`(*nlines*, *ncols*)  
`newterm`(*type*, *fd*)  
`newwin`(*lines*, *cols*, *begin\_y*, *begin\_x*)

`nl`()\*  
`nocbreak`()\*  
`nodelay`(*win*, *bf*)  
`noecho`()\*  
`nonl`()\*  
`noraw`()\*  
`overlay`(*win1*, *win2*)  
`overwrite`(*win1*, *win2*)  
`pnoutrefresh`(*pad*, *pminrow*, *pmincol*, *sminrow*,  
*smincol*, *smaxrow*, *smaxcol*)

`prefresh`(*pad*, *pminrow*, *pmincol*, *sminrow*,  
*smincol*, *smaxrow*, *smaxcol*)

`printw`(*fmt*, *arg1*, *arg2*, ...)

`raw`()\*  
`refresh`()\*  
`resetterm`()\*  
`resetty`()\*  
`saveterm`()\*  
`savetty`()\*  
`scanw`(*fmt*, *arg1*, *arg2*, ...)

`scroll`(*win*)  
`scrollok`(*win*, *flag*)  
`set_term`(*new*)  
`setscrreg`(*t*, *b*)  
`setterm`(*type*)  
`setupterm`(*term*, *filenum*, *errret*)  
`standend`()\*  
`standout`()\*  
`subwin`(*win*, *lines*, *cols*, *begin\_y*, *begin\_x*)

`touchwin`(*win*)  
`traceoff`()  
`traceon`()  
`typeahead`(*fd*)  
`unctrl`(*ch*)\*  
`waddch`(*win*, *ch*)  
`waddstr`(*win*, *str*)  
`wattroff`(*win*, *attrs*)  
`wattron`(*win*, *attrs*)  
`wattrset`(*win*, *attrs*)  
`wclear`(*win*)  
`wclrtoobot`(*win*)  
`wclrtoeol`(*win*)  
`wdelch`(*win*, *c*)  
`wdeleteln`(*win*)  
`werase`(*win*)  
`wgetch`(*win*)

create a new pad with given dimensions  
set up new terminal of given type to output on

create a new window  
set newline mapping  
unset cbreak mode  
enable nodelay input mode through getch  
unset echo mode  
unset newline mapping  
unset raw mode  
overlay *win1* on *win2*  
overwrite *win1* on top of *win2*

like `prefresh` but with no output until `doupdate` called

refresh from *pad* starting with given upper left  
corner of *pad* with output to given  
portion of screen

`printf` on *stdscr*  
set raw mode  
make current screen look like *stdscr*  
set tty modes to "out of curses" state  
reset tty flags to stored value  
save current modes as "in curses" state  
store current tty flags

`scanf` through *stdscr*  
`scroll` *win* one line  
allow terminal to scroll if *flag* != 0  
`now` talk to terminal *new*  
set user scrolling region to lines *t* through *b*  
establish terminal with given type

clear standout mode attribute  
set standout mode attribute

create a subwindow  
change all of *win*  
turn off debugging trace output  
turn on debugging trace output  
use file descriptor *fd* to check typeahead  
printable version of *ch*  
add char to *win*  
add string to *win*  
turn off *attrs* in *win*  
turn on *attrs* in *win*  
set *attrs* in *win* to *attrs*  
clear *win*  
clear to bottom of *win*  
clear to end of line on *win*  
delete char from *win*  
delete line from *win*  
erase *win*  
get a char through *win*



wgetstr(win, str)  
 winch(win)  
 winsch(win, c)  
 winsertln(win)  
 wmove(win, y, x)  
 wnoutrefresh(win)  
 wprintw(win, fmt, arg1, arg2, ...)

wrefresh(win)  
 wscanw(win, fmt, arg1, arg2, ...)

wsetscreg(win, t, b)  
 wstandend(win)  
 wstandout(win)

get a string through *win*  
 get char at current (y, x) in *win*  
 insert char into *win*  
 insert line into *win*  
 set current (y, x) co-ordinates on *win*  
 refresh but no screen output

printf on *win*  
 make screen look like *win*

scanf through *win*  
 set scrolling region of *win*  
 clear standout attribute in *win*  
 set standout attribute in *win*

## TERMINFO LEVEL ROUTINES

These routines should be called by programs wishing to deal directly with the terminfo database. Due to the low level of this interface, it is discouraged. Initially, *setupterm* should be called. This will define the set of terminal dependent variables defined in *terminfo(4)*. The include files *< curses.h >* and *< term.h >* should be included to get the definitions for these strings, numbers, and flags. Parameterized strings should be passed through *tparm* to instantiate them. All terminfo strings (including the output of *tparm*) should be printed with *tputs* or *putp*. Before exiting, *resetterm* should be called to restore the tty modes. (Programs desiring shell escapes or suspending with control Z can call *resetterm* before the shell is called and *fixterm* after returning from the shell.)

fixterm()  
 resetterm()  
 setupterm(term, fd, rc)

restore tty modes for terminfo use  
 (called by *setupterm*)  
 reset tty modes to state before program entry  
 read in database. Terminal type is the character string *term*, all output is to UNIX System file descriptor *fd*. A status value is returned in the integer pointed to by *rc*: 1 is normal. The simplest call would be *setupterm(0, 1, 0)* which uses all defaults.

tparm(str, p1, p2, ..., p9)

tputs(str, affcnt, putc)

instantiate string *str* with parms *p<sub>i</sub>*  
 apply padding info to string *str*.  
*affcnt* is the number of lines affected, or 1 if not applicable. *putc* is a putchar-like function to which the characters are passed, one at a time.

putp(str)

handy function that calls *tputs*  
 (str, 1, putchar)

vidputs(attrs, putc)

output the string to put terminal in video attribute mode *attrs*, which is any combination of the attributes listed below. Chars are passed to putchar-like function *putc*.

vidattr(attrs)

Like *vidputs* but outputs through *putc*

## TERMCAP COMPATIBILITY ROUTINES

These routines were included as a conversion aid for programs that use *termcap*. Their parameters are the same as for *termcap*. They are emulated using the *terminfo* database. They may go away at a later date.

tgetent(bp, name)  
 tgetflag(id)

look up *termcap* entry for name  
 get Boolean entry for id

## CURSES(3X)

tgetnum(id)	get numeric entry for id
tgetstr(id, area)	get string entry for id
tgoto(cap, col, row)	apply parms to given cap
tputs(cap, aff'cnt, fn)	apply padding to cap calling fn as putchar

### ATTRIBUTES

The following video attributes can be passed to the functions *attron*, *attroff*, *attrset*.

A_STANDOUT	Terminal's best highlighting mode
A_UNDERLINE	Underlining
A_REVERSE	Reverse video
A_BLINK	Blinking
A_DIM	Half bright
A_BOLD	Extra bright or bold
A_BLANK	Blanking (invisible)
A_PROTECT	Protected
A_ALTCHARSET	Alternate character set

### FUNCTION KEYS

The following function keys might be returned by *getch* if *keypad* has been enabled. Note that not all of these are currently supported, due to lack of definitions in *terminfo* or the terminal not transmitting a unique code when the key is pressed.

<i>Name</i>	<i>Value</i>	<i>Key name</i>
KEY_BREAK	0401	break key (unreliable)
KEY_DOWN	0402	The four arrow keys ...
KEY_UP	0403	
KEY_LEFT	0404	
KEY_RIGHT	0405	...
KEY_HOME	0406	Home key (upward+left arrow)
KEY_BACKSPACE	0407	backspace (unreliable)
KEY_F0	0410	Function keys. Space for 64 is reserved.
KEY_F(n)	(KEY_F0+(n))	Formula for fn.
KEY_DL	0510	Delete line
KEY_IL	0511	Insert line
KEY_DC	0512	Delete character
KEY_IC	0513	Insert char or enter insert mode
KEY_EIC	0514	Exit insert char mode
KEY_CLEAR	0515	Clear screen
KEY_EOS	0516	Clear to end of screen
KEY_EOL	0517	Clear to end of line
KEY_SF	0520	Scroll 1 line forward
KEY_SR	0521	Scroll 1 line backward (reverse)
KEY_NPAGE	0522	Next page
KEY_PPAGE	0523	Previous page
KEY_STAB	0524	Set tab
KEY_CTAB	0525	Clear tab
KEY_CATAB	0526	Clear all tabs
KEY_ENTER	0527	Enter or send (unreliable)
KEY_SRESET	0530	soft (partial) reset (unreliable)
KEY_RESET	0531	reset or hard reset (unreliable)
KEY_PRINT	0532	print or copy
KEY_LL	0533	home down or bottom (lower left)

WARNING

The plotting library *plot(3X)* and the curses library *curses(3X)* both use the names *erase()* and *move()*. The curses versions are macros. If you need both libraries, put the *plot(3X)* code in a different source file than the *curses(3X)* code, and/or `#undef move()` and `erase()` in the *plot(3X)* code.

S-3X

## LDAHREAD(3X)

### NAME

ldahread — read the archive header of a member of an archive file

### SYNOPSIS

```
#include <stdio.h>
#include <ar.h>
#include <filehdr.h>
#include <ldfcn.h>
```

```
int ldahread (ldptr, arhead)
LDFILE *ldptr;
ARCHDR *arhead;
```

### DESCRIPTION

If *TYPE(ldptr)* is the archive file magic number, *ldahread* reads the archive header of the common object file currently associated with *ldptr* into the area of memory beginning at *arhead*.

*Ldahread* returns **SUCCESS** or **FAILURE**. *Ldahread* will fail if *TYPE(ldptr)* does not represent an archive file, or if it cannot read the archive header.

The program must be loaded with the object file access routine library **libld.a**.

### SEE ALSO

ldclose(3X), ldopen(3X), ldfcn(4), ar(4).

## NAME

*ldclose*, *ldaclose* - close a common object file

## SYNOPSIS

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

int ldclose (ldptr)
LDFILE *ldptr;

int ldaclose (ldptr)
LDFILE *ldptr;
```

## DESCRIPTION

*Ldopen(3X)* and *ldclose* are designed to provide uniform access to both simple object files and object files that are members of archive files. Thus an archive of common object files can be processed as if it were a series of simple common object files.

If **TYPE**(*ldptr*) is not **ARTYPE**, *ldclose* will close the file and free the memory allocated to the **LDFILE** structure associated with *ldptr*. If **TYPE**(*ldptr*) is **ARTYPE**, and if there are any more files in the archive, *ldclose* will reinitialize **OFFSET**(*ldptr*) to the file address of the next archive member and return **FAILURE**. The **LDFILE** structure is prepared for a subsequent *ldopen*. In all other cases, *ldclose* returns **SUCCESS**.

*Ldaclose* closes the file and frees the memory allocated to the **LDFILE** structure associated with *ldptr* regardless of the value of **TYPE**(*ldptr*). *Ldaclose* always returns **SUCCESS**. The function is often used in conjunction with *ldaopen*.

The program must be loaded with the object file access routine library **libld.a**.

## SEE ALSO

*fclose(3S)*, *ldopen(3X)*, *ldfcn(4)*,

## LDFHREAD(3X)

### NAME

ldfhread - read the file header of a common object file

### SYNOPSIS

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

int ldfhread (ldptr, filehead)
LDFILE *ldptr;
FILHDR *filehead;
```

### DESCRIPTION

*Ldfhread* reads the file header of the common object file currently associated with *ldptr* into the area of memory beginning at *filehead*.

*Ldfhread* returns **SUCCESS** or **FAILURE**. *Ldfhread* will fail if it cannot read the file header.

In most cases the use of *ldfhread* can be avoided by using the macro **HEADER**(*ldptr*) defined in *ldfcn.h* [see *ldfcn(4)*]. The information in any field, *fieldname*, of the file header may be accessed using **HEADER**(*ldptr*),*fieldname*.

The program must be loaded with the object file access routine library **libl.a**.

### SEE ALSO

ldclose(3X), ldopen(3X), ldfcn(4).

## NAME

`ldgetname` - retrieve symbol name for a common object file symbol table entry

## SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <syms.h>
#include <ldfcn.h>

char *ldgetname (ldptr, symbol)
LDFILE *ldptr;
SYMENT *symbol;
```

## DESCRIPTION

*Ldgetname* returns a pointer to the name associated with *symbol* as a string. The string is contained in a static buffer local to *ldgetname* that is overwritten by each call to *ldgetname*, and therefore must be copied by the caller if the name is to be saved.

*Ldgetname* will return NULL (defined in *stdio.h*) for an object file if the name cannot be retrieved. This situation can occur:

- if the "string table" cannot be found,
- if not enough memory can be allocated for the string table,
- if the string table appears not to be a string table (for example, if an auxiliary entry is handed to *ldgetname* that looks like a reference to a name in a nonexistent string table), or
- if the name's offset into the string table is past the end of the string table.

Typically, *ldgetname* will be called immediately after a successful call to *ldtbread* to retrieve the name associated with the symbol table entry filled by *ldtbread*.

The program must be loaded with the object file access routine library **libld.a**.

## SEE ALSO

`ldclose(3X)`, `ldopen(3X)`, `ldtbseek(3X)`, `ldtbread(3X)`, `ldfcn(4)`.

## LDLREAD(3X)

### NAME

ldlread, ldlimit, ldlimitem - manipulate line number entries of a common object file function

### SYNOPSIS

```
#include <stdio.h>
#include <filehdr.h>
#include <linenum.h>
#include <ldfcn.h>

int ldlread(ldptr, fcndindx, linenum, linent)
LDFILE *ldptr;
long fcndindx;
unsigned short linenum;
LINENO linent;

int ldlimit(ldptr, fcndindx)
LDFILE *ldptr;
long fcndindx;

int ldlimitem(ldptr, linenum, linent)
LDFILE *ldptr;
unsigned short linenum;
LINENO linent;
```

### DESCRIPTION

*Ldlread* searches the line number entries of the common object file currently associated with *ldptr*. *Ldlread* begins its search with the line number entry for the beginning of a function and confines its search to the line numbers associated with a single function. The function is identified by *fcndindx*, the index of its entry in the object file symbol table. *Ldlread* reads the entry with the smallest line number equal to or greater than *linenum* into *linent*.

*Ldlimit* and *ldlimitem* together perform exactly the same function as *ldlread*. After an initial call to *ldlread* or *ldlimit*, *ldlimitem* may be used to retrieve a series of line number entries associated with a single function. *Ldlimit* simply locates the line number entries for the function identified by *fcndindx*. *Ldlimitem* finds and reads the entry with the smallest line number equal to or greater than *linenum* into *linent*.

*Ldlread*, *ldlimit*, and *ldlimitem* each return either **SUCCESS** or **FAILURE**. *Ldlread* will fail if there are no line number entries in the object file, if *fcndindx* does not index a function entry in the symbol table, or if it finds no line number equal to or greater than *linenum*. *Ldlimit* will fail if there are no line number entries in the object file or if *fcndindx* does not index a function entry in the symbol table. *Ldlimitem* will fail if it finds no line number equal to or greater than *linenum*.

The programs must be loaded with the object file access routine library **libl.a**.

### SEE ALSO

ldclose(3X), ldopen(3X), ldtbindex(3X), ldfcn(4).



## NAME

ldlseek, ldnlseek — seek to line number entries of a section of a common object file

## SYNOPSIS

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

```
int ldlseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;
```

```
int ldnlseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

## DESCRIPTION

*Ldlseek* seeks to the line number entries of the section specified by *sectindx* of the common object file currently associated with *ldptr*.

*Ldnlseek* seeks to the line number entries of the section specified by *sectname*.

*Ldlseek* and *ldnlseek* return SUCCESS or FAILURE. *Ldlseek* will fail if *sectindx* is greater than the number of sections in the object file; *ldnlseek* will fail if there is no section name corresponding with *\*sectname*. Either function will fail if the specified section has no line number entries or if it cannot seek to the specified line number entries.

Note that the first section has an index of *one*.

The program must be loaded with the object file access routine library **libld.a**.

## SEE ALSO

ldclose(3X), ldopen(3X), ldshread(3X), ldfcn(4).

## LDOHSEEK(3X)

### NAME

ldohseek — seek to the optional file header of a common object file

### SYNOPSIS

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

int ldohseek (ldptr)
LDFILE *ldptr;
```

### DESCRIPTION

*Ldohseek* seeks to the optional file header of the common object file currently associated with *ldptr*.

*Ldohseek* returns **SUCCESS** or **FAILURE**. *Ldohseek* will fail if the object file has no optional header or if it cannot seek to the optional header.

The program must be loaded with the object file access routine library **libld.a**.

### SEE ALSO

ldclose(3X), ldopen(3X), ldhread(3X), ldfcn(4).

## NAME

ldopen, ldaopen — open a common object file for reading

## SYNOPSIS

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

LDFILE *ldopen (filename, ldptr)
char *filename;
LDFILE *ldptr;

LDFILE *ldaopen (filename, oldptr)
char *filename;
LDFILE *oldptr;
```

## DESCRIPTION

*Ldopen* and *ldclose* (3X) are designed to provide uniform access to both simple object files and object files that are members of archive files. Thus an archive of common object files can be processed as if it were a series of simple common object files.

If *ldptr* has the value NULL, then *ldopen* will open *filename* and allocate and initialize the LDFILE structure, and return a pointer to the structure to the calling program.

If *ldptr* is valid and if *TYPE(ldptr)* is the archive magic number, *ldopen* will reinitialize the LDFILE structure for the next archive member of *filename*.

*Ldopen* and *ldclose* (3X) are designed to work in concert. *Ldclose* will return FAILURE only when *TYPE(ldptr)* is the archive magic number and there is another file in the archive to be processed. Only then should *ldopen* be called with the current value of *ldptr*. In all other cases, in particular whenever a new *filename* is opened, *ldopen* should be called with a NULL *ldptr* argument.

The following is a prototype for the use of *ldopen* and *ldclose* (3X).

```
/* for each filename to be processed */
ldptr = NULL;
do
{
    if ( (ldptr = ldopen(filename, ldptr)) != NULL )
    {
        /* check magic number */
        /* process the file */
    }
} while (ldclose(ldptr) == FAILURE);
```

If the value of *oldptr* is not NULL, *ldaopen* will open *filename* anew and allocate and initialize a new LDFILE structure, copying the *TYPE*, *OFFSET*, and *HEADER* fields from *oldptr*. *Ldaopen* returns a pointer to the new LDFILE structure. This new pointer is independent of the old pointer, *oldptr*. The two pointers may be used concurrently to read separate parts of the object file. For example, one pointer may be used to step sequentially through the relocation information, while the other is used to read indexed symbol table entries.

Both *ldopen* and *ldaopen* open *filename* for reading. Both functions return NULL if *filename* cannot be opened, or if memory for the LDFILE structure cannot be allocated. A successful open does not insure that the given file is a common object file or an archived object file.

The program must be loaded with the object file access routine library *libld.a*.

## SEE ALSO

fopen(3S), ldclose(3X), ldfcn(4)

## LDRSEEK(3X)

### NAME

`ldrseek`, `ldnrseek` — seek to relocation entries of a section of a common object file

### SYNOPSIS

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

int ldrseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnrseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

### DESCRIPTION

*Ldrseek* seeks to the relocation entries of the section specified by *sectindx* of the common object file currently associated with *ldptr*.

*Ldnrseek* seeks to the relocation entries of the section specified by *sectname*.

*Ldrseek* and *ldnrseek* return **SUCCESS** or **FAILURE**. *Ldrseek* will fail if *sectindx* is greater than the number of sections in the object file; *ldnrseek* will fail if there is no section name corresponding with *sectname*. Either function will fail if the specified section has no relocation entries or if it cannot seek to the specified relocation entries.

Note that the first section has an index of *one*.

The program must be loaded with the object file access routine library **libld.a**.

### SEE ALSO

`ldclose(3X)`, `ldopen(3X)`, `ldhread(3X)`, `ldfcn(4)`.

**NAME**

ldshread, ldnsread — read an indexed/named section header of a common object file

**SYNOPSIS**

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

```
#include <filehdr.h>
```

```
#include <scnhdr.h>
```

```
#include <ldfcn.h>
```

```
int ldshread (ldptr, sectindx, secthead)
```

```
LDFILE *ldptr;
```

```
unsigned short sectindx;
```

```
SCNHDR *secthead;
```

```
int ldnsread (ldptr, sectname, secthead)
```

```
LDFILE *ldptr;
```

```
char *sectname;
```

```
SCNHDR *secthead;
```

**DESCRIPTION**

*Ldshread* reads the section header specified by *sectindx* of the common object file currently associated with *ldptr* into the area of memory beginning at *secthead*.

*Ldnsread* reads the section header specified by *sectname* into the area of memory beginning at *secthead*.

*Ldshread* and *ldnsread* return SUCCESS or FAILURE. *Ldshread* will fail if *sectindx* is greater than the number of sections in the object file; *ldnsread* will fail if there is no section name corresponding with *sectname*. Either function will fail if it cannot read the specified section header.

Note that the first section header has an index of *one*.

The program must be loaded with the object file access routine library *libld.a*.

**SEE ALSO**

ldclose(3X), ldopen(3X), ldfcn(4).

## LDSSEEK(3X)

### NAME

ldsseek, ldnsseek — seek to an indexed/named section of a common object file

### SYNOPSIS

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

int ldsseek (ldptr, sectindx)
LDFILE *ldptr;
unsigned short sectindx;

int ldnsseek (ldptr, sectname)
LDFILE *ldptr;
char *sectname;
```

### DESCRIPTION

*Ldsseek* seeks to the section specified by *sectindx* of the common object file currently associated with *ldptr*.

*Ldnsseek* seeks to the section specified by *sectname*.

*Ldsseek* and *Ldnsseek* return SUCCESS or FAILURE. *Ldsseek* will fail if *sectindx* is greater than the number of sections in the object file; *Ldnsseek* will fail if there is no section name corresponding with *sectname*. Either function will fail if there is no section data for the specified section or if it cannot seek to the specified section.

Note that the first section has an index of *one*.

The program must be loaded with the object file access routine library **libld.a**.

### SEE ALSO

ldclose(3X), ldopen(3X), ldshread(3X), ldfcn(4).

## NAME

ldtindex — compute the index of a symbol table entry of a common object file

## SYNOPSIS

```
#include <stdio.h>
#include <filebdr.h>
#include <syms.h>
#include <ldfcn.h>

long ldtindex (ldptr)
LDFILE *ldptr;
```

## DESCRIPTION

*Ldtindex* returns the (**long**) index of the symbol table entry at the current position of the common object file associated with *ldptr*.

The index returned by *ldtindex* may be used in subsequent calls to *ldtbread*(3X). However, since *ldtindex* returns the index of the symbol table entry that begins at the current position of the object file, if *ldtindex* is called immediately after a particular symbol table entry has been read, it will return the index of the next entry.

*Ldtindex* will fail if there are no symbols in the object file, or if the object file is not positioned at the beginning of a symbol table entry.

Note that the first symbol in the symbol table has an index of *zero*.

The program must be loaded with the object file access routine library **libld.a**.

## SEE ALSO

ldclose(3X), ldopen(3X), ldtbread(3X), ldtbseek(3X), ldfcn(4).

## LDTBREAD(3X)

### NAME

ldtread — read an indexed symbol table entry of a common object file

### SYNOPSIS

```
#include <stdio.h>
#include <filehdr.b>
#include <syms.h>
#include <ldfcn.b>

int ldtread (ldptr, symindex, symbol)
LDFILE *ldptr;
long symindex;
SYMENT *symbol;
```

### DESCRIPTION

*Ldtread* reads the symbol table entry specified by **symindex** of the common object file currently associated with **ldptr** into the area of memory beginning at **symbol**.

*Ldtread* returns **SUCCESS** or **FAILURE**. *Ldtread* will fail if **symindex** is greater than the number of symbols in the object file, or if it cannot read the specified symbol table entry.

Note that the first symbol in the symbol table has an index of *zero*.

The program must be loaded with the object file access routine library **libld.a**.

### SEE ALSO

ldclose(3X), ldopen(3X), ldtbseek(3X), ldgetname(3X), ldfcn(4).



**NAME**

ldtbseek — seek to the symbol table of a common object file

**SYNOPSIS**

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

int ldtbseek (ldptr)
LDFILE *ldptr;
```

**DESCRIPTION**

*Ldtbseek* seeks to the symbol table of the object file currently associated with *ldptr*.

*Ldtbseek* returns **SUCCESS** or **FAILURE**. *Ldtbseek* will fail if the symbol table has been stripped from the object file, or if it cannot seek to the symbol table.

The program must be loaded with the object file access routine library **libld.a**.

**SEE ALSO**

ldclose(3X), ldopen(3X), ldtbread(3X), ldfcn(4).

# LOGNAME(3X)

## NAME

logname — return login name of user

## SYNOPSIS

char \*logname( )

## DESCRIPTION

*Logname* returns a pointer to the null-terminated login name; it extracts the \$LOGNAME variable from the user's environment.

This routine is kept in **/lib/libPW.a**.

## FILES

/etc/profile

## SEE ALSO

profile(4), environ(5),  
env(1), login(1) in the Runtime System manual.

## BUGS

The return values point to static data whose content is overwritten by each call. This method of determining a login name is subject to forgery.

## NAME

`malloc`, `free`, `realloc`, `calloc`, `malloc`, `mallinfo` — fast main memory allocator

## SYNOPSIS

```
#include <malloc.h>

char *malloc (size)
    unsigned size;

void free (ptr)
    char *ptr;

char *realloc (ptr, size)
    char *ptr;
    unsigned size;

char *calloc (nelem, elsize)
    unsigned nelem, elsize;

int malloc (cmd, value)
    int cmd, value;

struct mallinfo mallinfo (max)
    int max;
```

## DESCRIPTION

*Malloc* and *free* provide a simple general-purpose memory allocation package, which runs considerably faster than the *malloc(3C)* package. It is found in the library “*malloc*”, and is loaded if the option “*-lmalloc*” is used with *cc(1)* or *ld(1)*.

*Malloc* returns a pointer to a block of at least *size* bytes suitably aligned for any use.

The argument to *free* is a pointer to a block previously allocated by *malloc*; after *free* is performed this space is made available for further allocation, and its contents have been destroyed (but see *malloc* below for a way to change this behavior).

Undefined results will occur if the space assigned by *malloc* is overrun or if some random number is handed to *free*.

*Realloc* changes the size of the block pointed to by *ptr* to *size* bytes and returns a pointer to the (possibly moved) block. The contents will be unchanged up to the lesser of the new and old sizes.

*Calloc* allocates space for an array of *nelem* elements of size *elsize*. The space is initialized to zeros.

*Malloc* provides for control over the allocation algorithm. The available values for *cmd* are:

- |          |                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| M_MXFAST | Set <i>maxfast</i> to <i>value</i> . The algorithm allocates all blocks below the size of <i>maxfast</i> in large groups and then does them out very quickly. The default value for <i>maxfast</i> is 0.                                                                                                                                                                                             |
| M_NLBLKS | Set <i>numlblks</i> to <i>value</i> . The above mentioned “large groups” each contain <i>numlblks</i> blocks. <i>Numlblks</i> must be greater than 0. The default value for <i>numlblks</i> is 100.                                                                                                                                                                                                  |
| M_GRAIN  | Set <i>grain</i> to <i>value</i> . The sizes of all blocks smaller than <i>maxfast</i> are considered to be rounded up to the nearest multiple of <i>grain</i> . <i>Grain</i> must be greater than 0. The default value of <i>grain</i> is the smallest number of bytes which will allow alignment of any data type. Value will be rounded up to a multiple of the default when <i>grain</i> is set. |

## MALLOC(3X)

**M\_KEEP** Preserve data in a freed block until the next *malloc*, *realloc*, or *calloc*. This option is provided only for compatibility with the old version of *malloc* and is not recommended.

These values are defined in the `<malloc.h>` header file.

*Mallopt* may be called repeatedly, but may not be called after the first small block is allocated.

*Mallinfo* provides instrumentation describing space usage. It returns the structure:

```
struct mallinfo {
    int arena;           /* total space in arena */
    int ordblks;        /* number of ordinary blocks */
    int smblks;         /* number of small blocks */
    int hblkhd;         /* space in holding block headers */
    int hblks;          /* number of holding blocks */
    int usmblks;        /* space in small blocks in use */
    int fsmblks;        /* space in free small blocks */
    int uordblks;       /* space in ordinary blocks in use */
    int fordblks;       /* space in free ordinary blocks */
    int keepcost;       /* space penalty if keep option */
                        /* is used */
}
```

This structure is defined in the `<malloc.h>` header file.

Each of the allocation routines returns a pointer to space suitably aligned (after possible pointer coercion) for storage of any type of object.

### SEE ALSO

`brk(2)`, `malloc(3C)`.

### DIAGNOSTICS

*Malloc*, *realloc* and *calloc* return a NULL pointer if there is not enough available memory. When *realloc* returns NULL, the block pointed to by *ptr* is left intact. If *mallopt* is called after any allocation or if *cmd* or *value* are invalid, non-zero is returned. Otherwise, it returns zero.

### WARNINGS

This package usually uses more data space than *malloc(3C)*.

The code size is also bigger than *malloc(3C)*.

Note that unlike *malloc(3C)*, this package does not preserve the contents of a block when it is freed, unless the **M\_KEEP** option of *mallopt* is used.

Undocumented features of *malloc(3C)* have not been duplicated.

## NAME

plot — graphics interface subroutines

## SYNOPSIS

```

openpl ()
erase ()
label (s)
char *s;
line (x1, y1, x2, y2)
int x1, y1, x2, y2;
circle (x, y, r)
int x, y, r;
arc (x, y, x0, y0, x1, y1)
int x, y, x0, y0, x1, y1;
move (x, y)
int x, y;
cont (x, y)
int x, y;
point (x, y)
int x, y;
linemod (s)
char *s;
space (x0, y0, x1, y1)
int x0, y0, x1, y1;
closepl ()

```

## DESCRIPTION

These subroutines generate graphic output in a relatively device-independent manner. *Space* must be used before any of these functions to declare the amount of space necessary. See *plot(4)*. *Openpl* must be used before any of the others to open the device for writing. *Closepl* flushes the output.

*Circle* draws a circle of radius *r* with center at the point *(x, y)*.

*Arc* draws an arc of a circle with center at the point *(x, y)* between the points *(x0, y0)* and *(x1, y1)*.

String arguments to *label* and *linemod* are terminated by nulls and do not contain new-lines.

See *plot(4)* for a description of the effect of the remaining functions.

The library files listed below provide several flavors of these routines.

## FILES

/usr/lib/libplot.a	produces output for <i>tplot(1G)</i> filters
/usr/lib/lib300.a	for DASI 300
/usr/lib/lib300s.a	for DASI 300s
/usr/lib/lib450.a	for DASI 450
/usr/lib/lib4014.a	for TEKTRONIX 4014

## WARNINGS

In order to compile a program containing these functions in *file.c* it is necessary to use "cc *file.c* -lplot".

In order to execute it, it is necessary to use "a.out | tplot".

The above routines use `<stdio.h>`, which causes them to increase the size of programs, not otherwise using standard I/O, more than might be expected.

## **PLOT(3X)**

### **SEE ALSO**

**plot(4).**

**graph(1G), stat(1G), tplot(1G)** in the Runtime System manual.

## NAME

regcmp, regex - compile and execute regular expression

## SYNOPSIS

```
char *regcmp (string1 [, string2, ...], (char *)0)
char *string1, *string2, ...;

char *regex (re, subject[, ret0, ...])
char *re, *subject, *ret0, ...;

extern char *__loc1;
```

## DESCRIPTION

*Regcmp* compiles a regular expression and returns a pointer to the compiled form. *Malloc(3C)* is used to create space for the vector. It is the user's responsibility to free unneeded space so allocated. A NULL return from *regcmp* indicates an incorrect argument. *Regcmp(1)* has been written to generally preclude the need for this routine at execution time.

*Regex* executes a compiled pattern against the subject string. Additional arguments are passed to receive values back. *Regex* returns NULL on failure or a pointer to the next unmatched character on success. A global character pointer *\_\_loc1* points to where the match began. *Regcmp* and *regex* were mostly borrowed from the editor, *ed(1)*; however, the syntax and semantics have been changed slightly. The following are the valid symbols and their associated meanings.

- [ ]\* . These symbols retain their current meaning.
- \$ Matches the end of the string; \n matches a new-line.
- Within brackets the minus means *through*. For example, [a-z] is equivalent to [abcd...xyz]. The - can appear as itself only if used as the first or last character. For example, the character class expression [ ]- matches the characters ] and -.
- + A regular expression followed by + means *one or more times*. For example, [0-9]+ is equivalent to [0-9]0-9]\*.
- {m} {m,} {m,u} Integer values enclosed in {} indicate the number of times the preceding regular expression is to be applied. The value *m* is the minimum number and *u* is a number, less than 256, which is the maximum. If only *m* is present (e.g., {m}), it indicates the exact number of times the regular expression is to be applied. The value {m,} is analogous to {m,infinity}. The plus (+) and star (\*) operations are equivalent to {1,} and {0,} respectively.
- (...)\$n The value of the enclosed regular expression is to be returned. The value will be stored in the (n+1)th argument following the subject argument. At most ten enclosed regular expressions are allowed. *Regex* makes its assignments unconditionally.
- (... ) Parentheses are used for grouping. An operator, e.g., \*, +, {}, can work on a single character or a regular expression enclosed in parentheses. For example, (a\*(cb+)\*)\$0.

By necessity, all the above defined symbols are special. They must, therefore, be escaped to be used as themselves.

S-3X

## REGCMP(3X)

### EXAMPLES

Example 1:

```
char *cursor, *newcursor, *ptr;
...
newcursor = regex((ptr = regcmp(" \n", 0)), cursor);
free(ptr);
```

This example will match a leading new-line in the subject string pointed at by cursor.

Example 2:

```
char ret0[9];
char *newcursor, *name;
...
name = regcmp("[A-Za-z][A-Za-z0-9_]{0,7}")$0", '(char *)0');
newcursor = regex(name, "123Testing321", ret0);
```

This example will match through the string "Testing3" and will return the address of the character after the last matched character (cursor+11). The string "Testing3" will be copied to the character array *ret0*.

Example 3:

```
#include "file.i"
char *string, *newcursor;
...
newcursor = regex(name, string);
```

This example applies a precompiled regular expression in **file.i** [see *regcmp(1)*] against *string*.

This routine is kept in **/lib/<module>/libPW.a**, where *module* is either small or large.

### SEE ALSO

*malloc(3C)*  
*ed(1)*, *regemp(1)* in the Runtime System manual.

### BUGS

The user program may run out of memory if *regcmp* is called iteratively without freeing the vectors no longer required. The following user-supplied replacement for *malloc(3C)* reuses the same vector, saving time and space:

```
/* user's program */
...
char *
malloc(n)
unsigned n;
{
    static char rebuf[512];
    return (n <= sizeof rebuf) ? rebuf : NULL;
}
```



**NAME**

sputl, sgetl — access long integer data in a machine-independent fashion.

**SYNOPSIS**

```
void sputl (value, buffer)
long value;
char *buffer;

long sgetl (buffer)
char *buffer;
```

**DESCRIPTION**

*Sputl* takes the four bytes of the long integer *value* and places them in memory starting at the address pointed to by *buffer*. The ordering of the bytes is the same across all machines.

*Sgetl* retrieves the four bytes in memory starting at the address pointed to by *buffer* and returns the long integer value in the byte ordering of the host machine.

The combination of *sputl* and *sgetl* provides a machine-independent way of storing long numeric data in a file in binary form without conversion to characters.

A program which uses these functions must be loaded with the object-file access routine library *libld.a*.

# VPRINTF(3X)

## NAME

`vprintf`, `vfprintf`, `vsprintf` – print formatted output of a `varargs` argument list

## SYNOPSIS

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

int vprintf (format, ap)
char *format;
va_list ap;

int fprintf (stream, format, ap)
FILE *stream;
char *format;
va_list ap;

int sprintf (s, format, ap)
char *s, *format;
va_list ap;
```

## DESCRIPTION

`vprintf`, `vfprintf`, and `vsprintf` are the same as `printf`, `fprintf`, and `sprintf` respectively, except that instead of being called with a variable number of arguments, they are called with an argument list as defined by `varargs(5)`.

## EXAMPLE

The following demonstrates how `vfprintf` could be used to write an error routine.

```
#include <stdio.h>
#include <varargs.h>
.
.
.
/*
 * error should be called like
 * error(function_name, format, arg1, arg2...);
 */
/*VARARGS*/
void
error(va_list)
/* Note that the function_name and format arguments cannot be
 * separately declared because of the definition of varargs.
 */
va_dcl
{
    va_list args;
    char *fmt;

    va_start(args);
    /* print out name of function causing error */
    (void)fprintf(stderr, "ERROR in %s: ", va_arg(args, char *));
    fmt = va_arg(args, char *);
    /* print out remainder of message */
    (void)vfprintf(fmt, args);
    va_end(args);
    (void)abort( );
}
}
```

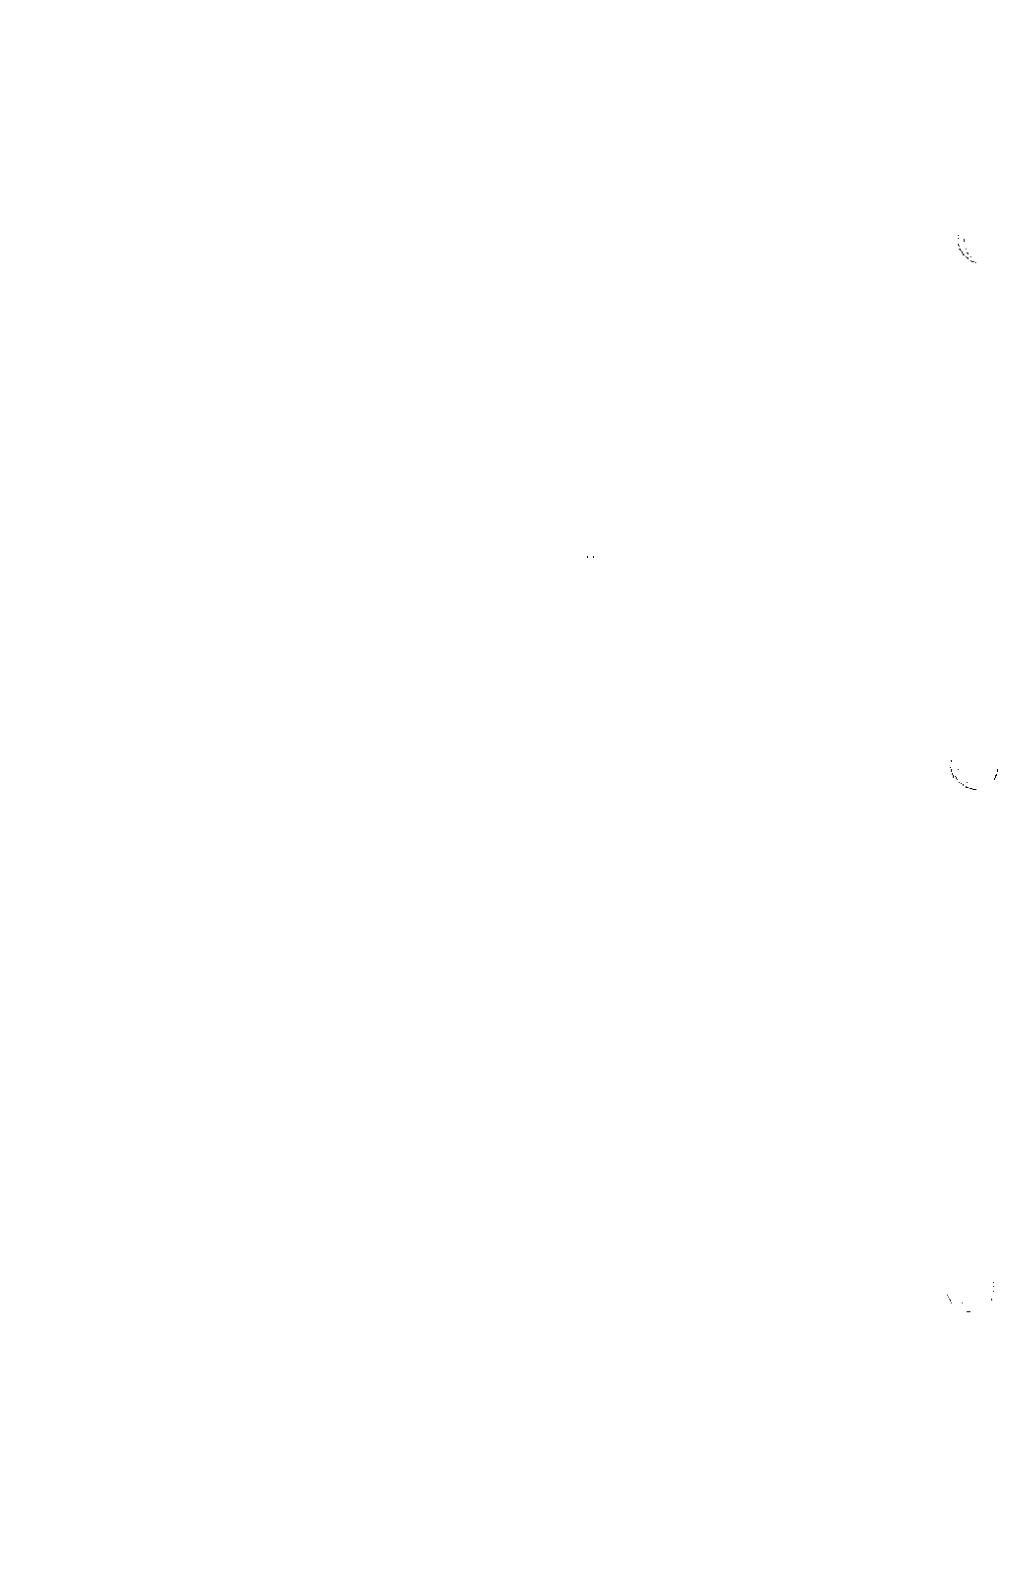
## SEE ALSO

`printf(3S)`, `varargs(5)`.

# TABLE OF CONTENTS OF FORTRAN LIBRARY FUNCTIONS

## 3F. Fortran Library Functions

<b>abort</b> .....	terminate FORTRAN program
<b>abs</b> .....	FORTRAN absolute value
<b>acos</b> .....	FORTRAN arccosine intrinsic function
<b>aimag</b> .....	FORTRAN imaginary part of complex argument
<b>aint</b> .....	FORTRAN integer part intrinsic function
<b>asin</b> .....	FORTRAN arcsine intrinsic function
<b>atan</b> .....	FORTRAN arctangent intrinsic function
<b>atan2</b> .....	FORTRAN arctangent intrinsic function
<b>bool</b> .....	FORTRAN bitwise Boolean functions
<b>conjg</b> .....	FORTRAN complex conjugate intrinsic function
<b>cos</b> .....	FORTRAN cosine intrinsic function
<b>cosh</b> .....	FORTRAN hyperbolic cosine intrinsic function
<b>dim</b> .....	positive difference intrinsic functions
<b>dprod</b> .....	double precision product intrinsic function
<b>exp</b> .....	FORTRAN exponential intrinsic function
<b>ftype</b> .....	explicit FORTRAN type conversion
<b>getarg</b> .....	return FORTRAN command-line argument
<b>getenv</b> .....	return FORTRAN environment variable
<b>iargc</b> .....	return number of command line arguments
<b>index</b> .....	return location of FORTRAN substring
<b>len</b> .....	return length of FORTRAN string
<b>log</b> .....	FORTRAN natural logarithm intrinsic function
<b>log10</b> .....	FORTRAN common logarithm intrinsic function
<b>max</b> .....	FORTRAN maximum-value functions
<b>mclock</b> .....	return FORTRAN time accounting
<b>min</b> .....	FORTRAN minimum-value functions
<b>mod</b> .....	FORTRAN remaindering intrinsic functions
<b>rand</b> .....	random number generator
<b>round</b> .....	FORTRAN nearest integer functions
<b>sign</b> .....	FORTRAN transfer-of-sign intrinsic function
<b>signal</b> .....	specify FORTRAN action on receipt of a system signal
<b>sin</b> .....	FORTRAN sine intrinsic function
<b>sinh</b> .....	FORTRAN hyperbolic sine intrinsic function
<b>sqrt</b> .....	FORTRAN square root intrinsic function
<b>strcmp</b> .....	string comparison intrinsic functions
<b>system</b> .....	issue a shell command from FORTRAN
<b>tan</b> .....	FORTRAN tangent intrinsic function
<b>tanh</b> .....	FORTRAN hyperbolic tangent intrinsic function



**NAME**

abort — terminate FORTRAN program

**SYNOPSIS**

call abort ( )

**DESCRIPTION**

*Abort* terminates the program which calls it, closing all open files truncated to the current position of the file pointer.

**DIAGNOSTICS**

When invoked, *abort* prints "FORTRAN abort routine called" on the standard error output.

**SEE ALSO**

abort(3C).

## ABS(3F)

### NAME

*abs*, *iabs*, *dabs*, *cabs*, *zabs* – FORTRAN absolute value

### SYNOPSIS

**integer** *i1*, *i2*  
**real** *r1*, *r2*  
**double precision** *dp1*, *dp2*  
**complex** *cx1*, *cx2*  
**double complex** *dx1*, *dx2*  
  
*r2* = *abs*(*r1*)  
  
*i2* = *iabs*(*i1*)  
*i2* = *abs*(*i1*)  
  
*dp2* = *dabs*(*dp1*)  
*dp2* = *abs*(*dp1*)  
  
*cx2* = *cabs*(*cx1*)  
*cx2* = *abs*(*cx1*)  
  
*dx2* = *zabs*(*dx1*)  
*dx2* = *abs*(*dx1*)

### DESCRIPTION

*Abs* is the family of absolute value functions. *Iabs* returns the integer absolute value of its integer argument. *Dabs* returns the double-precision absolute value of its double-precision argument. *Cabs* returns the complex absolute value of its complex argument. *Zabs* returns the double-complex absolute value of its double-complex argument. The generic form *abs* returns the type of its argument.

### SEE ALSO

*floor*(3M).

**NAME**

*acos*, *dacos* — FORTRAN arccosine intrinsic function

**SYNOPSIS**

real *r1*, *r2*

double precision *dp1*, *dp2*

*r2* = *acos*(*r1*)

*dp2* = *dacos*(*dp1*)

*dp2* = *acos*(*dp1*)

**DESCRIPTION**

*Acos* returns the real arccosine of its real argument. *Dacos* returns the double-precision arccosine of its double-precision argument. The generic form *acos* may be used with impunity as its argument will determine the type of the returned value.

**SEE ALSO**

*trig*(3M).

## AIMAG(3F)

### NAME

aimag, dimag — FORTRAN imaginary part of complex argument

### SYNOPSIS

real r

complex cxr

double precision dp

double complex cxd

r = aimag(cxr)

dp = dimag(cxd)

### DESCRIPTION

*Aimag* returns the imaginary part of its single-precision complex argument.

*Dimag* returns the double-precision imaginary part of its double-complex argument.



**NAME**

aint, dint — FORTRAN integer part intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
r2 = aint(r1)
dp2 = dint(dp1)
dp2 = aint(dp1)
```

**DESCRIPTION**

*Aint* returns the truncated value of its real argument in a real. *Dint* returns the truncated value of its double-precision argument as a double-precision value. *Aint* may be used as a generic function name, returning either a real or double-precision value depending on the type of its argument.

## ASIN(3F)

### NAME

asin, dasin — FORTRAN arcsine intrinsic function

### SYNOPSIS

**real** r1, r2

**double precision** dp1, dp2

**r2 = asin(r1)**

**dp2 = dasin(dp1)**

**dp2 = asin(dp1)**

### DESCRIPTION

*Asin* returns the real arcsine of its real argument. *Dasin* returns the double-precision arcsine of its double-precision argument. The generic form *asin* may be used with impunity as it derives its type from that of its argument.

### SEE ALSO

trig(3M).

**NAME**

*atan*, *datan* — FORTRAN arctangent intrinsic function

**SYNOPSIS**

real *r1*, *r2*  
double precision *dp1*, *dp2*  
*r2* = *atan*(*r1*)  
*dp2* = *datan*(*dp1*)  
*dp2* = *atan*(*dp1*)

**DESCRIPTION**

*Atan* returns the real arctangent of its real argument. *Datan* returns the double-precision arctangent of its double-precision argument. The generic form *atan* may be used with a double-precision argument returning a double-precision value.

**SEE ALSO**

*trig*(3M).

## ATAN2(3F)

### NAME

atan2, datan2 — FORTRAN arctangent intrinsic function

### SYNOPSIS

**real** r1, r2, r3

**double precision** dp1, dp2, dp3

r3 = atan2(r1, r2)

dp3 = datan2(dp1, dp2)

dp3 = atan2(dp1, dp2)

### DESCRIPTION

*Atan2* returns the arctangent of *arg1/arg2* as a real value. *Datan2* returns the double-precision arctangent of its double-precision arguments. The generic form *atan2* may be used with impunity with double-precision arguments.

### SEE ALSO

trig(3M).

**NAME**

and, or, xor, not, lshift, rshift — FORTRAN bitwise Boolean functions

**SYNOPSIS**

```
integer i, j, k
real a, b, c
double precision dp1, dp2, dp3

k = and(i, j)
c = or(a, b)
j = xor(i, a)
j = not(i)
k = lshift(i, j)
k = rshift(i, j)
```

**DESCRIPTION**

The generic intrinsic Boolean functions *and*, *or* and *xor* return the value of the binary operations on their arguments. *Not* is a unary operator returning the one's complement of its argument. *Lshift* and *rshift* return the value of the first argument shifted left or right, respectively, the number of times specified by the second (integer) argument.

The Boolean functions are generic, that is, they are defined for all data types as arguments and return values. Where required, the compiler will generate appropriate type conversions.

**NOTE**

Although defined for all data types, use of Boolean functions on any but integer data is bizarre and will probably result in unexpected consequences.

**BUGS**

The implementation of the shift functions may cause large shift values to deliver weird results.

## CONJG(3F)

### NAME

*conjg*, *dconjg* – FORTRAN complex conjugate intrinsic function

### SYNOPSIS

**complex** *cx1*, *cx2*  
**double complex** *dx1*, *dx2*  
*cx2* = *conjg*(*cx1*)  
*dx2* = *dconjg*(*dx1*)

### DESCRIPTION

*Conjg* returns the complex conjugate of its complex argument. *Dconjg* returns the double-complex conjugate of its double-complex argument.

**NAME**

*cos*, *dcos*, *ccos* — FORTRAN cosine intrinsic function

**SYNOPSIS**

real *r1*, *r2*  
double precision *dp1*, *dp2*  
complex *cx1*, *cx2*  
*r2* = *cos*(*r1*)  
*dp2* = *dcos*(*dp1*)  
*dp2* = *cos*(*dp1*)  
*cx2* = *ccos*(*cx1*)  
*cx2* = *cos*(*cx1*)

**DESCRIPTION**

*Cos* returns the real cosine of its real argument. *Dcos* returns the double-precision cosine of its double-precision argument. *Ccos* returns the complex cosine of its complex argument. The generic form *cos* may be used with impunity as its returned type is determined by that of its argument.

**SEE ALSO**

trig(3M).

## COSH(3F)

### NAME

cosh, dcosh — FORTRAN hyperbolic cosine intrinsic function

### SYNOPSIS

```
real r1, r2
double precision dp1, dp2
r2 = cosh(r1)
dp2 = dcosh(dp1)
dp2 = cosh(dp1)
```

### DESCRIPTION

*Cosh* returns the real hyperbolic cosine of its real argument. *Dcosh* returns the double-precision hyperbolic cosine of its double-precision argument. The generic form *cosh* may be used to return the hyperbolic cosine in the type of its argument.

### SEE ALSO

sinh(3M).



## NAME

dim, ddim, idim — positive difference intrinsic functions

## SYNOPSIS

```
integer a1,a2,a3
a3 = idim(a1,a2)
```

```
real a1,a2,a3
a3 = dim(a1,a2)
```

```
double precision a1,a2,a3
a3 = ddim(a1,a2)
```

## DESCRIPTION

These functions return:

```
  a1-a2  if a1 > a2
  0       if a1 <= a2
```

## **DPROD(3F)**

### **NAME**

dprod — double precision product intrinsic function

### **SYNOPSIS**

```
real a1,a2  
double precision a3  
a3 = dprod (a1,a2)
```

### **DESCRIPTION**

Dprod returns the double precision product of its real arguments.

**NAME**

*exp*, *dexp*, *cexp* — FORTRAN exponential intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
complex cx1, cx2

r2 = exp(r1)
dp2 = dexp(dp1)
dp2 = exp(dp1)
cx2 = clog(cx1)
cx2 = exp(cx1)
```

**DESCRIPTION**

*Exp* returns the real exponential function  $e^x$  of its real argument. *Dexp* returns the double-precision exponential function of its double-precision argument. *Cexp* returns the complex exponential function of its complex argument. The generic function *exp* becomes a call to *dexp* or *cexp* as required, depending on the type of its argument.

**SEE ALSO**

*exp*(3M).

## FTYPE(3F)

### NAME

int, ifix, idint, real, float, sngl, dble, cmplx, dcmplx, ichar, char — explicit FORTRAN type conversion

### SYNOPSIS

integer i, j  
real r, s  
double precision dp, dq  
complex cx  
double complex dcx  
character\*/ cb  
  
i = int(r)  
i = int(dp)  
i = int(cx)  
i = int(dcx)  
i = ifix(r)  
i = idint(dp)  
  
r = real(i)  
r = real(dp)  
r = real(cx)  
r = real(dcx)  
r = float(i)  
r = sngl(dp)  
  
dp = dble(i)  
dp = dble(r)  
dp = dble(cx)  
dp = dble(dcx)  
  
cx = cmplx(i)  
cx = cmplx(i, j)  
cx = cmplx(r)  
cx = cmplx(r, s)  
cx = cmplx(dp)  
cx = cmplx(dp, dq)  
cx = cmplx(dcx)  
  
dcx = dcmplx(i)  
dcx = dcmplx(i, j)  
dcx = dcmplx(r)  
dcx = dcmplx(r, s)  
dcx = dcmplx(dp)  
dcx = dcmplx(dp, dq)  
dcx = dcmplx(dcx)  
  
i = ichar(cb)  
ch = char(i)

### DESCRIPTION

These functions perform conversion from one data type to another.

The function **int** converts to *integer* form its *real*, *double precision*, *complex*, or *double complex* argument. If the argument is *real* or *double precision*, **int** returns the integer whose magnitude is the largest integer that does not exceed the magnitude of the argument and whose sign is the same as the sign of the argument (i.e., truncation). For complex types, the above rule is applied to the real part. **ifix** and **idint** convert only *real* and *double precision* arguments respectively.

The function **real** converts to *real* form an *integer*, *double precision*, *complex*, or *double complex* argument. If the argument is *double precision* or *double complex*, as much precision is kept as is possible. If the argument is one of the complex types, the real part is returned. **float** and **sngl** convert only *integer* and *double precision* arguments respectively.

The function **dblr** converts any *integer*, *real*, *complex*, or *double complex* argument to *double precision* form. If the argument is of a complex type, the real part is returned.

The function **cmplx** converts its *integer*, *real*, *double precision*, or *double complex* argument(s) to *complex* form.

The function **dcmplx** converts to *double complex* form its *integer*, *real*, *double precision*, or *complex* argument(s).

Either one or two arguments may be supplied to **cmplx** and **dcmplx**. If there is only one argument, it is taken as the real part of the complex type and an imaginary part of zero is supplied. If two arguments are supplied, the first is taken as the real part and the second as the imaginary part.

The function **ichar** converts from a character to an integer depending on the character's position in the collating sequence.

The function **char** returns the character in the *i*th position in the processor collating sequence where *i* is the supplied argument.

For a processor capable of representing *n* characters,

**ichar(char(i))** = *i* for  $0 \leq i < n$ , and

**char(ichar(ch))** = *ch* for any representable character *ch*.

## GETARG(3F)

### NAME

*getarg* — return FORTRAN command-line argument

### SYNOPSIS

**character\****N* *c*

**integer** *i*

*getarg*(*i*, *c*)

### DESCRIPTION

*Getarg* returns the *i*-th command-line argument of the current process. Thus, if a program were invoked via

foo arg1 arg2 arg3

*getarg*(2, *c*) would return the string "arg2" in the character variable *c*.

### SEE ALSO

*getopt*(3C).

**NAME**

*getenv* — return FORTRAN environment variable

**SYNOPSIS**

**character\*N c**

***getenv*( TMPDIR , c)**

**DESCRIPTION**

*Getenv* returns the character-string value of the environment variable represented by its first argument into the character variable of its second argument. If no such environment variable exists, all blanks will be returned.

**SEE ALSO**

*getenv*(3C), *environ*(5).

## IARGC(3F)

### NAME

`iargc`

### SYNOPSIS

```
integer i
i = iargc()
```

### DESCRIPTION

The *iargc* function returns the number of command line arguments passed to the program. Thus, if a program were invoked via

```
foo arg1 arg2 arg3
```

`iargc()` would return "3".

### SEE ALSO

`getarg(3F)`.



## NAME

*index* — return location of FORTRAN substring

## SYNOPSIS

character•N1 *ch1*

character•N2 *ch2*

integer *i*

*i* = *index* (*ch1*, *ch2*)

## DESCRIPTION

*Index* returns the location of substring *ch2* in string *ch1*. The value returned is the position at which substring *ch2* starts, or 0 if it is not present in string *ch1*.

## LEN(3F)

### NAME

len — return length of FORTRAN string

### SYNOPSIS

**character\*N ch**

**integer i**

**i = len(ch)**

### DESCRIPTION

*Len* returns the length of string *ch*.

**NAME**

**log, alog, dlog, clog** – FORTRAN natural logarithm intrinsic function

**SYNOPSIS**

**real r1, r2**  
**double precision dp1, dp2**  
**complex cx1, cx2**  
**r2 = alog(r1)**  
**r2 = log(r1)**  
**dp2 = dlog(dp1)**  
**dp2 = log(dp1)**  
**cx2 = clog(cx1)**  
**cx2 = log(cx1)**

**DESCRIPTION**

*Alog* returns the real natural logarithm of its real argument. *Dlog* returns the double-precision natural logarithm of its double-precision argument. *Clog* returns the complex logarithm of its complex argument. The generic function *log* becomes a call to *alog*, *dlog*, or *clog* depending on the type of its argument.

**SEE ALSO**

**exp(3M).**

## LOG10(3F)

### NAME

log10,alog10,dlog10 — FORTRAN common logarithm intrinsic function

### SYNOPSIS

```
real r1, r2
double precision dp1, dp2
r2 = alog10(r1)
r2 = log10(r1)
dp2 = dlog10(dp1)
dp2 = log10(dp1)
```

### DESCRIPTION

*Alog10* returns the real common logarithm of its real argument. *Dlog10* returns the double-precision common logarithm of its double-precision argument. The generic function *log10* becomes a call to *alog10* or *dlog10* depending on the type of its argument.

### SEE ALSO

*exp(3M)*.

## NAME

max, max0, amax0, max1, amax1, dmax1 — FORTRAN maximum-value functions

## SYNOPSIS

```
integer i, j, k, l
real a, b, c, d
double precision dp1, dp2, dp3

l = max(i, j, k)
c = max(a, b)
dp = max(a, b, c)
k = max0(l, j)
a = amax0(i, j, k)
i = max1(a, b)
d = amax1(a, b, c)
dp3 = dmax1(dp1, dp2)
```

## DESCRIPTION

The maximum-value functions return the largest of their arguments (of which there may be any number). *Max* is the generic form which can be used for all data types and takes its return type from that of its arguments (which must all be of the same type). *Max0* returns the integer form of the maximum value of its integer arguments; *amax0*, the real form of its integer arguments; *max1*, the integer form of its real arguments; *amax1*, the real form of its real arguments; and *dmax1*, the double-precision form of its double-precision arguments.

## SEE ALSO

min(3F).

## MCLOCK(3F)

### NAME

mclock — return FORTRAN time accounting

### SYNOPSIS

integer i

i = mclock( )

### DESCRIPTION

*Mclock* returns time accounting information about the current process and its child processes. The value returned is the sum of the current process's user time and the user and system times of all child processes.

### SEE ALSO

times(2), clock(3C), system(3F).

## NAME

min, min0, amin0, min1, amin1, dmin1 — FORTRAN minimum-value functions

## SYNOPSIS

integer i, j, k, l  
 real a, b, c, d  
 double precision dp1, dp2, dp3

l = min(i, j, k)  
 c = min(a, b)  
 dp = min(a, b, c)  
 k = min0(i, j)  
 a = amin0(i, j, k)  
 i = min1(a, b)  
 d = amin1(a, b, c)  
 dp3 = dmin1(dp1, dp2)

## DESCRIPTION

The minimum-value functions return the minimum of their arguments (of which there may be any number). *Min* is the generic form which can be used for all data types and takes its return type from that of its arguments (which must all be of the same type). *Min0* returns the integer form of the minimum value of its integer arguments; *amin0*, the real form of its integer arguments; *min1*, the integer form of its real arguments; *amin1*, the real form of its real arguments; and *dmin1*, the double-precision form of its double-precision arguments.

## SEE ALSO

max(3F).

## MOD(3F)

### NAME

`mod`, `amod`, `dmod` – FORTRAN remaindering intrinsic functions

### SYNOPSIS

integer `i`, `j`, `k`  
real `r1`, `r2`, `r3`  
double precision `dp1`, `dp2`, `dp3`  
`k` = `mod`(`i`, `j`)  
`r3` = `amod`(`r1`, `r2`)  
`r3` = `mod`(`r1`, `r2`)  
`dp3` = `dmod`(`dp1`, `dp2`)  
`dp3` = `mod`(`dp1`, `dp2`)

### DESCRIPTION

*Mod* returns the integer remainder of its first argument divided by its second argument. *Amod* and *dmod* return, respectively, the real and double-precision whole number remainder of the integer division of their two arguments. The generic version *mod* will return the data type of its arguments.



**NAME**

*irand*, *rand*, *srand* — random number generator

**SYNOPSIS**

call *srand*(*iseed*)

*i* = *irand*()

*x* = *rand*( )

**DESCRIPTION**

*Irand* generates successive pseudo-random numbers in the range from 0 to  $2^{**}15-1$ . *Rand* generates pseudo-random numbers distributed in (0, 1.0). *Srand* uses its integer argument to re-initialize the seed for successive invocations of *irand* and *rand*.

**SEE ALSO**

*rand*(3C).

## ROUND(3F)

### NAME

anint, dnint, nint, idnint — FORTRAN nearest integer functions

### SYNOPSIS

**integer i**  
**real r1, r2**  
**double precision dp1, dp2**  
**r2 = anint(r1)**  
**i = nint(r1)**  
**dp2 = anint(dp1)**  
**dp2 = dnint(dp1)**  
**i = nint(dp1)**  
**i = idnint(dp1)**

### DESCRIPTION

*Anint* returns the nearest whole real number to its real argument [i.e.,  $\text{int}(a+0.5)$  if  $a \geq 0$ ,  $\text{int}(a-0.5)$  otherwise]. *Dnint* does the same for its double-precision argument. *Nint* returns the nearest integer to its real argument. *Idnint* is the double-precision version. *Anint* is the generic form of *anint* and *dnint*, performing the same operation and returning the data type of its argument. *Nint* is also the generic form of *idnint*.

**NAME**

*sign*, *isign*, *dsign* – FORTRAN transfer-of-sign intrinsic function

**SYNOPSIS**

**integer** *i*, *j*, *k*  
**real** *r1*, *r2*, *r3*  
**double precision** *dp1*, *dp2*, *dp3*  
*k* = *isign*(*i*, *j*)  
*k* = *sign*(*i*, *j*)  
*r3* = *sign*(*r1*, *r2*)  
*dp3* = *dsign*(*dp1*, *dp2*)  
*dp3* = *sign*(*dp1*, *dp2*)

**DESCRIPTION**

*Isign* returns the magnitude of its first argument with the sign of its second argument. *Sign* and *dsign* are its real and double-precision counterparts, respectively. The generic version is *sign* and will devolve to the appropriate type depending on its arguments.

## **SIGNAL(3F)**

### **NAME**

signal — specify FORTRAN action on receipt of a system signal

### **SYNOPSIS**

**integer i**  
**external integer intfnc**  
**call signal(i, intfnc)**

### **DESCRIPTION**

*Signal* allows a process to specify a function to be invoked upon receipt of a specific signal. The first argument specifies which fault or exception; the second argument the specific function to be invoked.

### **SEE ALSO**

kill(2), signal(2).

## NAME

*sin*, *dsin*, *csin* — FORTRAN sine intrinsic function

## SYNOPSIS

real *r1*, *r2*  
double precision *dp1*, *dp2*  
complex *cx1*, *cx2*  
*r2* = *sin*(*r1*)  
*dp2* = *dsin*(*dp1*)  
*dp2* = *sin*(*dp1*)  
*cx2* = *csin*(*cx1*)  
*cx2* = *sin*(*cx1*)

## DESCRIPTION

*Sin* returns the real sine of its real argument. *Dsin* returns the double-precision sine of its double-precision argument. *Csin* returns the complex sine of its complex argument. The generic *sin* function becomes *dsin* or *csin* as required by argument type.

## SEE ALSO

trig(3M).

## SINH(3F)

### NAME

`sinh`, `dsinh` – FORTRAN hyperbolic sine intrinsic function

### SYNOPSIS

```
real r1, r2
double precision dp1, dp2
r2 = sinh(r1)
dp2 = dsinh(dp1)
dp2 = sinh(dp1)
```

### DESCRIPTION

*Sinh* returns the real hyperbolic sine of its real argument. *Dsinh* returns the double-precision hyperbolic sine of its double-precision argument. The generic form *sinh* may be used to return a double-precision value when given a double-precision argument.

### SEE ALSO

`sinh(3M)`.

## NAME

*sqrt*, *dsqrt*, *csqrt* — FORTRAN square root intrinsic function

## SYNOPSIS

real *r1*, *r2*  
double precision *dpl*, *dp2*  
complex *cx1*, *cx2*  
  
*r2* = *sqrt*(*r1*)  
  
*dp2* = *dsqrt*(*dpl*)  
*dp2* = *sqrt*(*dpl*)  
  
*cx2* = *csqrt*(*cx1*)  
*cx2* = *sqrt*(*cx1*)

## DESCRIPTION

*Sqrt* returns the real square root of its real argument. *Dsqrt* returns the double-precision square root of its double-precision argument. *Csqrt* returns the complex square root of its complex argument. *Sqrt*, the generic form, will become *dsqrt* or *csqrt* as required by its argument type.

## SEE ALSO

*exp*(3M).

## STRCMP(3F)

### NAME

lge, lgt, lle, llt — string comparison intrinsic functions

### SYNOPSIS

character\*N    a1, a2  
logical l

l = lge (a1,a2)

l = lgt (a1,a2)

l = lle (a1,a2)

l = llt (a1,a2)

### DESCRIPTION

These functions return .TRUE. if the inequality holds and .FALSE. otherwise.



**NAME**

system — issue a shell command from FORTRAN

**SYNOPSIS**

character•N c

call system(c)

**DESCRIPTION**

*System* causes its character argument to be given to *sh*(1) as input, as if the string had been typed at a terminal. The current process waits until the shell has completed.

**SEE ALSO**

exec(2), system(3S).

sh(1) in the Runtime System manual.

## TAN(3F)

### NAME

tan, dtan — FORTRAN tangent intrinsic function

### SYNOPSIS

real r1, r2

double precision dp1, dp2

r2 = tan(r1)

dp2 = dtan(dp1)

dp2 = tan(dp1)

### DESCRIPTION

*Tan* returns the real tangent of its real argument. *Dtan* returns the double-precision tangent of its double-precision argument. The generic *tan* function becomes *dtan* as required with a double-precision argument.

### SEE ALSO

trig(3M).

**NAME**

*tanh*, *dtanh* — FORTRAN hyperbolic tangent intrinsic function

**SYNOPSIS**

```
real r1, r2
double precision dp1, dp2
r2 = tanh(r1)
dp2 = dtanh(dp1)
dp2 = tanh(dp1)
```

**DESCRIPTION**

*Tanh* returns the real hyperbolic tangent of its real argument. *Dtanh* returns the double-precision hyperbolic tangent of its double-precision argument. The generic form *tanh* may be used to return a double-precision value given a double-precision argument.

**SEE ALSO**

*sinh*(3M).



# TABLE OF CONTENTS OF FILE FORMATS

## 4. File Formats

intro.....	introduction to file formats
a.out.....	assembler and link editor output
acct.....	per-process accounting file format
ar.....	common archive file format
checklist.....	list of file systems processed by fsck
core.....	format of core image file
cpio.....	format of cpio archive
dialinfo.....	dial procedure data base
dir.....	format of directories
errfile.....	error-log file format
filehdr.....	file header for object files
fs.....	format of system volume
fspec.....	format specification in text files
gettydefs.....	speed and terminal settings used by getty
gps.....	graphical primitive string, format of graphical files
group.....	group file
inittab.....	script for the init process
inode.....	format of an i-node
issue.....	issue identification file
ldfcn.....	object file access routines
linenum.....	line number entries in a common object file
master.....	master device information table
mmtab.....	mounted file system table
passwd.....	password file
plot.....	graphics interface
prch.....	file format for card images
profile.....	setting up an environment at login time
reloc.....	relocation information for an object file
scsf file.....	format of SCCS file
scnhdr.....	section header for an object file
syms.....	common object file symbol table format
term.....	format of compiled term file
terminfo.....	terminal capability data base
unrp.....	utmp and wtmp entry formats



**NAME**

intro - introduction to file formats

**DESCRIPTION**

This section outlines the formats of various files. The C **struct** declarations for the file formats are given where applicable. Usually, these structures can be found in the directories **/usr/include** or **/usr/include/sys**.

References of the type *name*(1M) refer to entries found in Section 1 of the Runtime System manual.

**This page intentionally left blank.**



## NAME

a.out - common assembler and link editor output

## DESCRIPTION

The file name **a.out** is the output file from the assembler *as*(1) and the link editor *ld*(1). Both programs will make *a.out* executable if there were no errors in assembling or linking and no unresolved external references.

A common object file consists of a file header, a UNIX system header, a table of section headers, relocation information, (optional) line numbers, a symbol table, and a string table. The order is given below.

```

File header.
UNIX system header.
Section 1 header.
...
Section n header.
Section 1 data.
...
Section n data.
Section 1 relocation.
...
Section n relocation.
Section 1 line numbers.
...
Section n line numbers.
Symbol table.
String table.

```

The last three parts of an object file (line numbers, symbol table and string table) may be missing if the program was linked with the *-s* option of *ld*(1) or if they were removed by *strip*(1). Also note that the relocation information will be absent if there were no unresolved external references after linking. The string table exists only if the symbol table contains symbols with names longer than eight characters.

The sizes of each section (contained in the header, discussed below) are in bytes and are even.

When an **a.out** file is loaded for execution, two or more memory segments are set up. The number of segments depends on the memory model selected at compile time and the program's size. The loader builds these segments from the data sections of the object file, according to the parameters in the headers.

In the small memory model, programs have one data segment (containing a stack and all of the data), and one code (program text) segment.

In the large memory model, programs have a separate stack segment and one or more segments for both data and code. Each segment in the iAPX 286 is limited to 64k bytes of information; so for these models, the compiler creates multiple code and data segments when necessary.

The data segment is extended during program execution as requested by the *brk*(2) system call. When necessary in large model programs, additional memory segments are allocated to accommodate such extensions.

The value of a word in the text or data portions that is not a reference to an undefined external symbol is exactly the value that will appear in memory when the file is executed. If a word in the text involves a reference to an undefined external symbol, the storage class of the symbol-table entry for that word will be marked as an "external symbol", and the

## A.OUT(4)

section number will be set to 0. When the file is processed by the link editor and the external symbol becomes defined, the value of the symbol will be added to the word in the file.

### File Header

The format of the filehdr header is

```
struct filehdr
{
    unsigned short f_magic; /* magic number */
    unsigned short f_nscns; /* number of sections */
    long          f_timdat; /* time and date stamp */
    long          f_symptr; /* file ptr to symtab */
    long          f_nsyms; /* # symtab entries */
    unsigned short f_opthdr; /* sizeof(opt hdr) */
    unsigned short f_flags; /* flags */
};
```

### UNIX System Header

The format of the UNIX system header is

```
typedef struct aouthdr
{
    short  magic; /* magic number */
    short  vstamp; /* version stamp */
    long   tsize; /* text size in bytes, padded */
    long   dsize; /* initialized data (.data) */
    long   bsize; /* uninitialized data (.bss) */
    long   entry; /* entry point */
    long   text_start; /* base of text used for this file */
    long   data_start; /* base of data used for this file */
} AOUTHDR;
```

### Section Header

The format of the section header is

```

struct scnhdr
{
    char          s_name[SYMNMLEN]; /* section name */
    long         s_paddr; /* physical address */
    long         s_vaddr; /* virtual address */
    long         s_size; /* section size */
    long         s_scnptr; /* file ptr to raw data */
    long         s_relptr; /* file ptr to relocation */
    long         s_lnnoptr; /* file ptr to line numbers */
    unsigned short s_nreloc; /* # reloc entries */
    unsigned short s_nlnno; /* # line number entries */
    long         s_flags; /* flags */
};

```

### Relocation

Object files have one relocation entry for each relocatable reference in the text or data. If relocation information is present, it will be in the following format:

```

struct reloc
{
    long         r_vaddr; /* (virtual) address of reference */
    long         r_symndx; /* index into symbol table */
    short        r_type; /* relocation type */
};

```

The start of the relocation information is *s\_relptr* from the section header. If there is no relocation information, *s\_relptr* is 0.

## Symbol Table

The format of each symbol in the symbol table is

```

#define SYMNMLEN 8
#define FILNMLEN 14
#define SYMESZ 18 /* the size of a SYMENT */

struct syment
{
    union /* all ways to get a symbol name */
    {
        char _n_name[SYMNMLEN]; /* name of symbol */
        struct
        {
            long _n_zeroes; /* == 0L if in string table */
            long _n_offset; /* location in string table */
        } _n_n;
        char *_n_nptr[2]; /* allows overlaying */
    } _n;
    unsigned long n_value; /* value of symbol */
    short n_scnm; /* section number */
    unsigned short n_type; /* type and derived type */
    char n_sclass; /* storage class */
    char n_numaux; /* number of aux entries */
};

#define n_name _n_n_name
#define n_zeroes _n_n_n_n_zeroes
#define n_offset _n_n_n_n_offset
#define n_nptr _n_n_nptr[1]

```

Some symbols require more information than a single entry; they are followed by *auxiliary entries* that are the same size as a symbol entry. The format follows.

```

union auxent {
    struct {
        long    x_tagndx;
        union {
            struct {
                unsigned short x_lno;
                unsigned short x_size;
            } x_lnsz;
            long    x_fsize;
        } x_misc;
        union {
            struct {
                long    x_lnoptr;
                long    x_endndx;
            } x_fcn;
            struct {
                unsigned short x_dimen[DIMNUM];
            } x_ary;
        } x_fenary;
        unsigned short x_tvndx;
    } x_sym;

    struct {
        char    x_fname[FILNMLEN];
    } x_file;

    struct {
        long    x_scnlen;
        unsigned short x_nreloc;
        unsigned short x_nlinno;
    } x_scn;

    struct {
        long    x_tvfill;
        unsigned short x_tvlen;
        unsigned short x_tvran[2];
    } x_tv;
};

```

Indexes of symbol table entries begin at *zero*. The start of the symbol table is *f\_symptr* (from the file header) bytes from the beginning of the file. If the symbol table is stripped, *f\_symptr* is 0. The string table (if one exists) begins at *f\_symptr* + (*f\_nsyms* \* SYMESZ) bytes from the beginning of the file.

**SEE ALSO**

brk(2), filehdr(4), ldfcn(4), linenum(4), reloc(4), scnhdr(4), syms(4), as(1), cc(1), ld(1) in the Runtime System manual.

## ACCT(4)

### NAME

acct - per-process accounting file format

### SYNOPSIS

```
#include <sys/acct.h>
```

### DESCRIPTION

Files produced as a result of calling *acct(2)* have records in the form defined by `<sys/acct.h>`, whose contents are:

```
typedef ushort comp_t; /* "floating point" */
                          /* 13-bit fraction, 3-bit exponent */

struct acct
{
    char    ac_flag;      /* Accounting flag */
    char    ac_stat;     /* Exit status */
    ushort  ac_uid;
    ushort  ac_gid;
    dev_t   ac_tty;
    time_t  ac_btime;    /* Beginning time */
    comp_t  ac_utime;    /* acctng user time in clock ticks */
    comp_t  ac_stime;    /* acctng system time in clock ticks */
    comp_t  ac_etime;    /* acctng elapsed time in clock ticks */
    comp_t  ac_mem;      /* memory usage in clicks */
    comp_t  ac_io;       /* chars trnsfrd by read/write */
    comp_t  ac_rw;       /* number of block reads/writes */
    char    ac_comm[8];  /* command name */
};

extern struct acct    acctbuf;
extern struct inode   *acctp; /* inode of accounting file */

#define AFORK 01        /* has executed fork, but no exec */
#define ASU   02        /* used super-user privileges */
#define ACCTF 0300     /* record type: 00 = acct */
```

In *ac\_flag*, the AFORK flag is turned on by each *fork(2)* and turned off by an *exec(2)*. The *ac\_comm* field is inherited from the parent process and is reset by any *exec*. Each time the system charges the process with a clock tick, it also adds to *ac\_mem* the current process size, computed as follows:

$$(\text{data size}) + (\text{text size}) / (\text{number of in-core processes using text})$$

The value of  $ac\_mem / (ac\_stime + ac\_utime)$  can be viewed as an approximation to the mean process size, as modified by text-sharing.

The structure `tacct.h`, which resides with the source files of the accounting commands, represents the total accounting format used by the various accounting commands:

```

/*
 * total accounting (for acct period), also for day
 */

struct tacct {
    uid_t      ta_uid;      /* userid */
    char       ta_name[8]; /* login name */
    float      ta_cpu[2];  /* cum. cpu time, p/np (mins) */
    float      ta_kcore[2]; /* cum kcore-minutes, p/np */
    float      ta_con[2];  /* cum. connect time, p/np, mins */
    float      ta_du;      /* cum. disk usage */
    long       ta_pc;      /* count of processes */
    unsigned short ta_sc;  /* count of login sessions */
    unsigned short ta_dc;  /* count of disk samples */
    unsigned short ta_fee; /* fee for special services */
};

```

#### SEE ALSO

`acct(2)`, `exec(2)`, `fork(2)`,  
`acct(1M)`, `acctcom(1)` in the Runtime System manual.

#### BUGS

The `ac_mem` value for a short-lived command gives little information about the actual size of the command, because `ac_mem` may be incremented while a different command (e.g., the shell) is being executed by the process.

## AR(4)

### NAME

ar — common archive file format

### DESCRIPTION

The archive command *ar(1)* is used to combine several files into one. Archives are used mainly as libraries to be searched by the link editor *ld(1)*.

Each archive begins with the archive magic string.

```
#define ARMAG "!<arch>\n" /* magic string */
#define SARMAG 8 /* length of magic string */
```

Each archive which contains common object files [see *a.out(4)*] includes an archive symbol table. This symbol table is used by the link editor *ld(1)* to determine which archive members must be loaded during the link edit process. The archive symbol table (if it exists) is always the first file in the archive (but is never listed) and is automatically created and/or updated by *ar*.

Following the archive magic string are the archive file members. Each file member is preceded by a file member header which is of the following format:

```
#define ARFN:AG "\n" /* header trailer string */

struct ar_hdr /* file member header */
(
    char ar_name[16]; /* '/' terminated file member name */
    char ar_date[12]; /* file member date */
    char ar_uid[6]; /* file member user identification */
    char ar_gid[6]; /* file member group identification */
    char ar_mode[8]; /* file member mode (octal) */
    char ar_size[10]; /* file member size */
    char ar_fm[2]; /* header trailer string */
);
```

All information in the file member headers is in printable ASCII. The numeric information contained in the headers is stored as decimal numbers (except for *ar\_mode* which is in octal). Thus, if the archive contains printable files, the archive itself is printable.

The *ar\_name* field is blank-padded and slash (/) terminated. The *ar\_date* field is the modification date of the file at the time of its insertion into the archive. Common format archives can be moved from system to system as long as the portable archive command *ar(1)* is used.

Each archive file member begins on an even byte boundary; a newline is inserted between files if necessary. Nevertheless the size given reflects the actual size of the file exclusive of padding.

Notice there is no provision for empty areas in an archive file.

If the archive symbol table exists, the first file in the archive has a zero length name (i.e., *ar\_name[0] == '/'*). The contents of this file are as follows:

- The number of symbols. Length: 4 bytes.
- The array of offsets into the archive file. Length: 4 bytes \* "the number of symbols".
- The name string table. Length: *ar\_size* - (4 bytes \* ("the number of symbols" + 1)).



The number of symbols and the array of offsets are managed with *sgetl* and *sputl*. The string table contains exactly as many null terminated strings as there are elements in the offsets array. Each offset from the array is associated with the corresponding name from the string table (in order). The names in the string table are all the defined global symbols found in the common object files in the archive. Each offset is the location of the archive header for the associated symbol.

**SEE ALSO**

*sput(3X)*, *a.out(4)*.

*ar(1)*, *ld(1)*, *strip(1)* in the Runtime System manual.

**CAVEATS**

*Strip(1)* will remove all archive symbol entries from the header. The archive symbol entries must be restored via the *ts* option of the *ar(1)* command before the archive can be used with the link editor *ld(1)*.

## CHECKLIST(4)

### NAME

checklist - list of file systems processed by fsck

### DESCRIPTION

*Checklist* resides in directory */etc* and contains a list of, at most, 15 *special file* names. Each *special file* name is contained on a separate line and corresponds to a file system. Each file system will then be automatically processed by the *fsck(1M)* command.

### SEE ALSO

*fsck(1M)* in the Runtime System manual.

**NAME**

core — format of core image file

**DESCRIPTION**

The UNIX system writes out a core image of a terminated process when any of various errors occur. See *signal(2)* for the list of reasons; the most common are memory violations, illegal instructions, bus errors, and user-generated quit signals. The core image is called **core** and is written in the process's working directory (provided it can be; normal access controls apply). A process with an effective user ID different from the real user ID will not produce a core image.

The first section of the core image is a copy of the system's per-user data for the process, including the registers as they were at the time of the fault. The size of this section depends on the parameter *usize*, which is defined in */usr/include/sys/param.h*. The second section of the core image is a copy of the process's local descriptor table(LDT). The size of this section is specified in the per-user data field, *u\_lsize*. The remainder represents the actual contents of the user's core area when the core image was written. If the text segment is read-only and shared, it is not dumped.

The format of the information in the first section is described by the *user* structure of the system, defined in */usr/include/sys/user.h*. The important stuff not detailed therein is the locations of the registers, which are outlined in */usr/include/sys/reg.h*.

The format of the information in the second section is described by the *seg\_desc* structure of the system, defined in */usr/include/sys/seg.h*. The LDT is an array of these structures.

**SEE ALSO**

*setuid(2)*, *signal(2)*,  
*crash(1M)*, *sdb(1)* in the Runtime System manual.



**NAME**

dialinfo - dial procedure data base

**SYNOPSIS**

/usr/lib/uucp/dialinfo

**DESCRIPTION**

The *dialinfo* file is a database of *dialers*. Each dialer is a procedure for connecting a terminal line to some kind of remote computer system. Most *dialer* entries dial out on a modem with built-in auto-dial capability. Some more esoteric entries resolve baud rate switching with a remote *getty*, perform auto-login, and even detect/defeat remote dialback modems.

Each *dialer* entry describes the operation of a state machine. This state machine is capable of sending characters to the remote, controlling the switchhook, sending break, adjusting baud-rate, and printing user messages on *stderr*. State transition decisions are made according to responses from the remote, response timeout, carrier detection, and retry count. The state machine may be partially or fully customized with user environment variables.

*Dialinfo* contains an extensive diagnostic capability both to support initial debugging, and to deal with day-to-day problems on real telephone lines and down systems.

The entry format of *dialinfo* is modeled after *terminfo*. Entries have the general form:

```
name1 | name2 | . . .
      key1=str1, key2=str2, . . .
      . . .
      keyn=strn, . . .
```

Where:

**name?** Descriptive dialer name. For an auto-dial modem this is usually manufacturer name and model number. For a dialer to be used, this name must appear in field 3 of a corresponding entry in the *L-devices* file.

**key?** Dialinfo keyword.

**str?** Keyword definition string. A definition may span several lines. Only commas need be escaped.

The name entry must begin in column 1. Keyword entries may immediately follow the first comma, or may be indented and placed on subsequent lines. Blank lines, and comment lines beginning with # are ignored.

Keywords include:

**delay** Modem string to get a 2 second dial delay.

**walt** Modem string to wait for dial tone or secondary dial tone.

**star** Modem string for the \* key on a touch tone phone.

**pound** Modem string for the # key on a touch tone phone.

**flash** Modem string to flash off hook for 1 second.

**retry** Initial retry count. Specifies the number of retries to be performed before announcing failure.

## DIALINFO(4)

- sn** Definition of machine state *n*. There are 100 possible states numbered 0 to 99.
- use** The specified string is a dialer entry from which remaining undefined keywords may be taken. Presently defined keywords are not affected. Any keyword definitions which follow in the current dialer entry are ignored.

An environment variable string may be placed anywhere in a *dialinfo* definition. The form  $\$(name)$  inserts the value of the environment variable *name*. The form  $\$(name-default)$  inserts the value of *name* if it exists, and otherwise literally inserts the *default* string.

Machine state definitions contain a sequence of commands performed in order from left to right. Possible commands include:

- B** Transmit 250 millisecond break signal.
- Cn** Set connect option. If *n* is 0, use 8 bits, no parity, L-devices baud rate and *clocal* mode for modem communication. If *n* is 1, switch parity, character size and baud rate for communication with the remote system; carrier detect is still disabled. If *n* is 2, proceed as with 1, but enable carrier detect if this is an ACU line or *call-modem* is non-zero. If *n* is 3, set options as with 1, but always enable carrier detect. If *n* is 4, set options as with 1, but enable carrier detect and wait for carrier to be established. If carrier is not seen within the timeout period set by the last S command, take action according to the last H command.
- Dn** Drop DTR (hang up line) for *n* seconds. When DTR is restored, baud rate, parity etc is set as if C0 was executed.
- E "string"** Write *string* to *stderr*. Presumably *stderr* is directed to a user terminal, or a log file.
- F "string"** Declare that *string* should be sent to *stderr* on any subsequent dial failure. Remains in effect until canceled.
- G s** Immediately transfer control to state *s*.
- Hs** Declare that the carrier lost error recovery state is *s*. If a subsequent read of the communication line fails because carrier sense is enabled and DCD is false, control will be transferred to state *s*.
- M "string"** Write *string* to the communication line.
- Nb** Change the communication line baud rate to speed *b*.
- Pn** Pause for *n* seconds.
- Rn** Decrement the retry count, and fail if the result is negative.
- Sn** Declare that the timeout used when waiting for carrier or a communication line response is *n* seconds.
- Ts** Declare that the timeout recovery state is *s*. Any subsequent timeout will cause a transfer to this state.

- Un* Execute (use) the text of state definition *n* as a subroutine. This command nests up to 10 deep.
- [*string*] *s* Declare that a transfer to state *s* should be performed when *string* is received from the remote.

The state parameter denoted as *s* in the above entries is one of the following:

- n* Transfer control to state *n*, where *n* is a decimal number in the range 0 to 99.
- +
- 
- n* Exit, returning the error code *-n*.

The following escapes are recognised in E, F, M, and [ ] command strings:

- \*nnn* The octal character *nnn*.
- ^*c* The control character derived by the logical and of the ASCII character *c* and the octal mask 037. This is the character transmitted by a standard ASCII keyboard when the control key is held down, and the character *c* is depressed.
- \*c* Standard C language escapes \b (backspace), \f (formfeed), \t (tab), \n (newline), and \r (return).
- %*n* Field number *n* of the current L-devices entry.
- %{*var*} The user environment variable *var*.
- %*an* The unmodified telephone number string.
- %*N* Converted telephone number described below.

The following characters have special meaning in the %*N* character strings, and are replaced with corresponding *dialinfo* strings as described below.

- \* or : (*star*) Dials the \* on a touch-tone telephone.
- # or ; (*pound*) Dials the # on a touch-tone telephone.
- (*pause*) Pauses 2 seconds.
- =orw (*wait*) Wait for secondary dial tone.
- f (*flash*) Flash offhook for one second.

## DIALER OPERATION

The state machine is initialized to "C0H- S10T-", and execution begins at state 0.

As each state is entered, any previous [ ] command strings are cleared. Command execution then proceeds in order from left to right.

If a command (eg G) is encountered that causes a change of state, remaining commands in the current state are not performed, and execution continues in the new state.

Otherwise, when all the commands in a state have been executed, a read operation is performed on the communication line. Incoming characters are then matched against previously declared [ ] command strings. If a match is found, the corresponding state transfer occurs. If carrier sense was enabled by a previous C command, and carrier is lost, the last H command is honored. If neither of these events occur within the timeout specified by the last S command, a timeout is detected. Action is then taken according to the last T command.

## DIALINFO(4)

### EXAMPLES

The following is a stripped-down definition for a Hayes Smartmodem.

```
hayes,  
star=*, pound=#, flash=H0\, H1,  
delay=\, , wait=\, , retry=2,  
s0=M"\rAT DT%N\r" S60 [CONNECT] +,
```

Below is a more complicated entry to handle a Vadic 3451.

```
vadic | va3451,  
delay=K, wait=KK, retry=5,  
s0=P1 M"^E\r" [*] 1 S2 T10,  
s1=P1 M"D\r" [NUMBER] 2,  
s2=P1 M"%N\r" P1 "\r" [DIALING] 3,  
s3=E"Dialing %n..." S30 H11 C4 [CONNECT] 4 T11,  
s4=E"Connected." G+,  
s10=F"No response from modem." R1 D1 G0,  
s11=F"Dial failed." R1 E"No Carrier, retrying..." D1 G0,
```

The next entry gets its complete definition from the DIALCUSTOM environment variable, and defaults to hayes if that variable is undefined.

```
custom,  
${DIALCUSTOM-usc=hayes} ,
```

### FILES

```
/usr/lib/uucp/dialinfo      Dial procedure data base.  
/usr/lib/uucp/L-devices    UUCP device file.
```

### AUTHOR

Gene H. Olson, Quest Research, Burnsville MN.

### SEE ALSO

```
dialer(1)           Modern dial-out program.  
dialprint(1)       Prints dialer entries.  
getty(1)           For information on dial-in lines.  
uucp(1)            UUCP interface information.  
dial(3)            C library dial procedure.  
terminfo(4)       Terminal capability data base.  
acu(7)             For System V phone number conventions.  
term(7)            Terminal device information.
```



**NAME**  
dir — format of directories

**SYNOPSIS**  
**#include** <sys/dir.h>

**DESCRIPTION**

A directory behaves exactly like an ordinary file, except that no user may write into a directory. The fact that a file is a directory is indicated by a bit in the flag word of its i-node entry [see *fs(4)*]. The structure of a directory entry as given in the include file is:

```
#ifndef DIRSIZ
#define DIRSIZ    14
#endif
struct direct
{
    ino_t  d_ino;
    char  d_name[DIRSIZ];
};
```

By convention, the first two entries in each directory are for . and .. The first is an entry for the directory itself. The second is for the parent directory. The meaning of .. is modified for the root directory of the master file system; there is no parent, so .. has the same meaning as ..

**SEE ALSO**  
*fs(4)*.

# ERRFILE(4)

## NAME

errfile — error-log file format

## DESCRIPTION

When hardware errors are detected by the system, an error record is generated and passed to the error-logging daemon for recording in the error log for later analysis. The default error log is `/usr/adm/errfile`.

The format of an error record depends on the type of error that was encountered. Every record, however, has a header with the following format:

```
struct errhdr {
    short      e_type;      /* record type */
    short      e_len;      /* bytes in record (inc hdr) */
    time_t     e_time;     /* time of day */
};
```

The permissible record types are as follows:

```
#define E_GOTS      010      /* start for UNIX System/TS */
#define E_GORT      011      /* start for UNIX system/RT */
#define E_STOP      012      /* stop */
#define E_TCHG      013      /* time change */
#define E_CCHG      014      /* configuration change */
#define E_BLK       020      /* block device error */
#define E_STRAY     030      /* stray interrupt */
#define E_PRTY      031      /* memory parity */
#define E_PIO       041      /* 3B20 computer programmed
                             * I/O */
#define E_IOP       042      /* 3B20 computer I/O
                             * processor */
```

Some records in the error file are of an administrative nature. These include the startup record that is entered into the file when logging is activated, the stop record that is written if the daemon is terminated “gracefully”, and the time-change record that is used to account for changes in the system’s time-of-day. These records have the following formats:

```
struct estart {
    short      e_cpu;      /* CPU type */
    struct utsname e_name; /* system names */
#ifdef u3b
    short      e_mmr3;     /* contents mem mgmt reg 3 */
    long       e_syssize;  /* 11/70 system memory size */
    short      e_bconf;    /* block dev configuration */
#endif
#ifdef u3b
    int        e_mmcnt;    /* kbytes per array */
#endif
};

#define eend errhdr /* record header */

struct etimchg {
    time_t     e_ntime;    /* new time */
};
```

Stray interrupts cause a record with the following format to be logged:

```
struct estray {
#ifdef u3b
    uint          e_saddr;    /* stray loc or device addr */
#else
    physadr      e_saddr;    /* stray loc or device addr */
    short        e_sbacty;   /* active block devices */
#endif
};
```

Memory subsystem errors on 3B20 computer and the iAPX 286 and VAX 11/70 processors cause the following record to be generated:

```
struct eparity {
#ifdef u3b
    int          e_parreg[3]; /* 3B computer memory
                             * registers */
#else
    short        e_parreg[4]; /* memory subsys registers */
#endif
};
```

Memory subsystem errors on VAX-11/780 processors cause the following record to be generated:

```
struct ememory {
    int          e_sbier;
    int          e_memcad;
};
```

Memory subsystem errors on Intel-iAPX 286 processors cause the following record to be generated:

```
struct ememory {
    ushort       e_esr;      /* ECC error status register*/
    ushort       e_hard;    /* flag to indicate a hard err*/
};
```

Error records for block devices have the following format:

```
struct eblock {
#ifdef u3b
    ushort       e_num;     /* device number */
    struct iostat {
        long      io_ops;   /* number read/writes */
        long      io_misc;  /* number "other" operations */
        ushort    io_unlog; /* number unlogged errors */
    }
    e_stats;
    short        e_bflags;  /* read/write, error, etc */
    daddr_t      e_bnum;    /* logical block number */
    uint         e_bytes;   /* number bytes to transfer */
    union ptbl {
        int page[64];      /* page table entries */
        union ptbl *pnxt;
    }
    e_ptbl;
    struct ptbl  e_ptbl;    /* page table for transfer */
    uint         e_voff;    /* offset into page table */
    uint         e_stat1;   /* status word 1 */
    uint         e_stat2;   /* status word 2 */
#endif
};
```

```

#ifndef u3b
    dev_t      e_dev;      /* "true" major + minor dev no */
    physadr    e_regloc;   /* controller address */
    short      e_bacty;    /* other block I/O activity */
    struct iostat {
        long    io_ops;    /* number read/writes */
        long    io_misc;   /* number "other" operations */
        ushort  io_unlog;  /* number unlogged errors */
    }
    short      e_stats;
    short      e_bflags;   /* read/write, error, etc */
#endif
#define iAPX286
    short      e_secoff;   /* logical dev start sector */
#else
    short      e_cyloff;   /* logical dev start cyl */
#endif
#define
    daddr_t    e_bnum;    /* logical block number */
    ushort     e_bytes;   /* number bytes to transfer */
    paddr_t    e_memadd;  /* buffer memory address */
    ushort     e_rtry;    /* number retries */
    short      e_nreg;    /* number device registers */
#endif
#ifdef vax
    struct mba_regs {
        long mba_csr;
        long mba_cr;
        long mba_sr;
        long mba_var;
        long mba_vcr;
    } e_mba;
#endif
};

```

The following values are used in the *e\_bflags* word:

```

#define E_WRITE  0      /* write operation */
#define E_READ   1      /* read operation */
#define E_NOIO   02     /* no I/O pending */
#define E_PHYS   04     /* physical I/O */
#define E_MAP    010    /* Unibus map in use */
#define E_ERROR  020    /* I/O failed */

```

For iAPX 286 processors, a configuration change message is sent to the error logging daemon whenever a block device driver is attached or detached.

```

struct econfchg {
    char    e_trudev;    /* "true" major device number */
    char    e_cflag;    /* driver attached or detached */
};
#define E_ATCH  1
#define E_DTCH  0

```

iAPX 286 processors use the following structure for system accounting:

```

struct iotime {
    struct iostat ios;
    long    io_bcmt;    /* total blocks transferred */
    time_t  io_act;    /* total controller active time */
    time_t  io_resp;   /* total block response time */
};
#define io_cnt ios.io_ops

```

The following error records are for the 3B20 computer only:

```

struct epio {
    char          e_chan;      /* programmed I/O (pio) error */
    char          e_dev;      /* which channel */
    uint          e_chstat;   /* which dev on channel */
    uint          e_cmd;      /* channel status */
}

struct eiop {
    char          e_unit;     /* I/O processor (iop) error */
    uint          e_word0;    /* unit number */
    uint          e_word1;   /* iop report word */
}

```

The "true" major device numbers that identify the failing device are as follows:

<i>Digital Equipment</i>	<i>AT&amp;T</i>	<i>Intel Equipment</i>
#define RK0 0	#define DFC0 0	#define WN0 0
#define RP0 1	#define IOPO 1	
#define RF0 2	#define MTO 2	
#define TM0 3		
#define TCO 4		
#define HP0 5		
#define HT0 6		
#define HS0 7		
#define RLO 8		
#define HP1 9		
#define HP2 10		
#define HP3 11		

SEE ALSO

errdemon(1M) in the Runtime System manual.

## FILEHDR(4)

### NAME

filehdr - file header for common object files

### SYNOPSIS

```
#include "filehdr.h"
```

### DESCRIPTION

Every common object file begins with a 20-byte header. The following C **struct** declaration is used:

```
struct filehdr
{
    unsigned short f_magic ; /* magic number */
    unsigned short f_nsens ; /* number of sections */
    long          f_tmdat ; /* time & date stamp */
    long          f_sympr ; /* file ptr to symtab */
    long          f_nsyms ; /* # symtab entries */
    unsigned short f_opthdr ; /* sizeof(opt hdr) */
    unsigned short f_flags ; /* flags */
};
```

*F\_sympr* is the byte offset into the file at which the symbol table can be found. Its value can be used as the offset in *fseek(3S)* to position an I/O stream to the symbol table. There are two possible magic numbers corresponding to the two possible addressing models available:

```
#define I286SMAGIC 0512
#define I286LMAGIC 0522
```

The value in *f\_tmdat* is obtained from the *time(2)* system call.

Flag bits currently defined are:

```
#define F_RELFLG      00001 /* relocation entries stripped */
#define F_EXEC       00002 /* file is executable */
#define F_LNNO       00004 /* line numbers stripped */
#define F_LSYMS      00010 /* local symbols stripped */
#define F_MINMAL     00020 /* minimal object file */
#define F_UPDATE     00040 /* update file, ogen produced */
#define F_SWABD      00100 /* file is "pre-swabbed" */
#define F_AR16WR     00200 /* 16-bit DEC host */
#define F_AR32WR     00400 /* 32-bit DEC host */
#define F_AR32W      01000 /* non-DEC host */
#define F_PATCH      02000 /* "patch" list in opt hdr */
#define F_80186      010000 /* need to run on an 80186 or 80286 */
#define F_80286      020000 /* need to run on an 80286 */
```

### SEE ALSO

*time(2)*, *fseek(3S)*, *a.out(4)*.

Flag bits currently defined are:

```
#define F_RELFLG 0x0001 /* relocation entries stripped */
#define F_EXEC 0x0002 /* file is executable */
#define F_LNNO 0x0004 /* line numbers stripped */
#define F_LSYMS 0x0008 /* local symbols stripped */
#define F_MINMAL 0x0010 /* minimal object file */
#define F_UPDATE 0x0020 /* update file, ogen produced */
#define F_SWABD 0x0040 /* file is "pre-swabbed" */
#define F_AR16WR 0x0080 /* 16-bit DEC host */
#define F_AR32WR 0x0100 /* 32-bit DEC host */
#define F_AR32W 0x0200 /* non-DEC host */
#define F_PATCH 0x0400 /* "patch" list in opt_hdr */
```

#### Map of COFF to STL

The following table shows the mapping of the (virtual format) COFF file-header fields to their actual STL locations. These conventions are used:

The fileheader is referred to as "filehdr" in COFF and "stlhdr" in STL.

The system header is referred to as "aouthdr" in COFF and "exthdr" (extended header) in STL.

The section headers are referred to as "scnhdr" in COFF and "sechdr" in STL.

#### Filehdr

```
filehdr->f_magic = stlhdr->fh_magic;
filehdr->f_nscns = total number of text, data,
                  and bss sections
filehdr->f_timdat = exthdr->e_timdat;
if (sechdr->se_type == S_SYMINFO)
    /* section header for symbol table */
    filehdr->f_sympr = sechdr->se_scnptr;
filehdr->f_nsyms = (unsigned short) (stlhdr->fh_syntblsiz /
                                     (long) SYMESZ);

filehdr->f_opthdr = stlhdr->fh_exthdrsz;
filehdr->f_flags = exthdr->e_flags;
```

#### Both STL and COFF

*f\_sympr* is the byte offset into the file at which the symbol table can be found. This value can be used as the offset in *fseek*(3S) to position an I/O stream to the symbol table. The COFF UNIX system header is 36 bytes on the 3B20 computer, 28 bytes otherwise. The valid magic number is:

```
#define I286MAGIC 0x0206 /* iAPX 286 processor */
```

The value in *f\_timdat* (*e\_timdat*) is obtained from the *time*(2) system call.

#### SEE ALSO

*time*(2), *fseek*(3S), *a.out*(4).

## FS(4)

### NAME

file system — format of system volume

### SYNOPSIS

```
#include <sys/filsys.h>
#include <sys/types.h>
#include <sys/param.h>
```

### DESCRIPTION

Every file system storage volume has a common format for certain vital information. Every such volume is divided into a certain number of 512-byte long sectors. Sector 0 is unused and is available to contain a bootstrap program or other information.

Sector 1 is the *super block*. The format of a super block is:

```
/*
 * Structure of the super block
 */
struct filsys
{
    ushort    s_ isize;           /* size in blocks of i-list */
    daddr_t   s_ fsize;          /* size in blocks of entire volume */
    short     s_ nfree;          /* number of addresses in s_ free */
    daddr_t   s_ free[NICFREE];  /* free block list */
    short     s_ ninode;         /* number of i-nodes in s_ inode */
    ino_t     s_ inode[NICINOD]; /* free i-node list */
    char      s_ flock;          /* lock during free list manipulation */
    char      s_ ilock;          /* lock during i-list manipulation */
    char      s_ fmod;           /* super block modified flag */
    char      s_ ronly;          /* mounted read-only flag */
    time_t    s_ time;           /* last super block update */
    short     s_ dinfo[4];       /* device information */
    daddr_t   s_ tfree;          /* total free blocks */
    ino_t     s_ tinode;         /* total free i-nodes */
    char      s_ fname[6];       /* file system name */
    char      s_ fpack[6];       /* file system pack name */
    long      s_ fill[13];       /* ADJUST to make size of filsys
    be 512 */

    long      s_ magic;          /* magic number to denote new
    file system */

    long      s_ type;           /* type of new file system */
};

#define FsMAGIC 0xfd187e20      /* s_ magic number */

#define Fs1b    1               /* 512-byte block */
#define Fs2b    2               /* 1024-byte block */
```

*S\_type* indicates the file system type. Currently, two types of file systems are supported: the original 512-byte oriented and the new improved 1024-byte oriented. *S\_magic* is used to distinguish the original 512-byte oriented file systems from the newer file systems. If this field is not equal to the magic number, *FsMAGIC*, the type is assumed to be *Fs1b*, otherwise the *s\_type* field is used. In the following description, a block is then determined by the type. For the original 512-byte oriented file system, a block is 512 bytes. For the 1024-byte oriented file system, a block is 1024 bytes or two sectors. The operating system takes care of all conversions from logical block numbers to physical sector numbers.



## NAME

file system – format of system volume

## SYNOPSIS

```
#include <sys/filsys.h>
#include <sys/types.h>
#include <sys/param.h>
```

## DESCRIPTION

Every file system storage volume has a common format for certain vital information. Every such volume is divided into a certain number of 512-byte long sectors. Sector 0 is unused and is available to contain a bootstrap program or other information.

Sector 1 is the *super-block*. The format of a super-block is:

```
/*
 * Structure of the super-block
 */
struct filsys
{
    ushort    s_ysize;           /* size in blocks of i-list */
    daddr_t   s_fsize;          /* size in blocks of entire volume */
    short     s_nfree;          /* number of addresses in s_free */
    daddr_t   s_free[NICFREE];  /* free block list */
    short     s_ninode;         /* number of i-nodes in s_inode */
    ino_t     s_inode[NICINOD]; /* free i-node list */
    char      s_flock;          /* lock during free list manipulation */
    char      s_ilock;          /* lock during i-list manipulation */
    char      s_fmod;           /* super block modified flag */
    char      s_ronly;          /* mounted read-only flag */
    time_t    s_time;           /* last super block update */
    short     s_dinfo[4];       /* device information */
    daddr_t   s_tfree;          /* total free blocks */
    ino_t     s_tinode;         /* total free i-nodes */
    char      s_fname[6];       /* file system name */
    char      s_fpack[6];       /* file system pack name */
    long      s_fill[14];        /* ADJUST to make size of filsys
                                be 512 */
    long      s_state;          /* file system state */
    long      s_magic;          /* magic number to denote new
                                file system */
    long      s_type;           /* type of new file system */
};

#define FsMAGIC 0xfd187e20      /* s_magic number */
#define Fs1b 1                  /* 512-byte block */
#define Fs2b 2                  /* 1024-byte block */
#define FsOKAY 0x7c269d38       /* s_state: clean */
#define FsACTIVE 0x5e72d81a     /* s_state: active */
#define FsBAD 0xc096f43         /* s_state: bad root */
```

*S\_type* indicates the file system type. Currently, two types of file systems are supported: the original 512-byte oriented and the new improved 1024-byte oriented. *S\_magic* is used to distinguish the original 512-byte oriented file systems from the newer file systems. If this field is not equal to the magic number, *FsMAGIC*, the type is assumed to be *Fs1b*, otherwise the *s\_type* field is used. In the following description, a block is then determined by the type. For the original 512-byte oriented file system, a block is 512 bytes. For the 1024-byte oriented file system, a block

is 1024 bytes or two sectors. The operating system takes care of all conversions from logical block numbers to physical sector numbers. *S\_state* indicates the state of the file system. A cleanly unmounted, not damaged file system is indicated by the *FsOKAY* state. After the file system has been mounted for update, the state is changed to *FsACTIVE*. A special case is used for the root file system. If the root file system appears damaged at boot time, it is mounted but marked *FsBAD*.

*S\_izise* is the address of the first data block after the i-list; the i-list starts just after the super-block, namely in block 2; thus the i-list is *s\_izise*-2 blocks long. *S\_fsize* is the first block not potentially available for allocation to a file. These numbers are used by the system to check for bad block numbers; if an "impossible" block number is allocated from the free list or is freed, a diagnostic is written on the on-line console. Moreover, the free array is cleared, so as to prevent further allocation from a presumably corrupted free list.

The free list for each volume is maintained as follows. The *s\_free* array contains, in *s\_free*[1], ..., *s\_free*[*s\_nfree*-1], up to 49 numbers of free blocks. *S\_free*[0] is the block number of the head of a chain of blocks constituting the free list. The first long in each free-chain block is the number (up to 50) of free-block numbers listed in the next 50 longs of this chain member. The first of these 50 blocks is the link to the next member of the chain. To allocate a block: decrement *s\_nfree*, and the new block is *s\_free*[*s\_nfree*]. If the new block number is 0, there are no blocks left, so give an error. If *s\_nfree* became 0, read in the block named by the new block number, replace *s\_nfree* by its first word, and copy the block numbers in the next 50 longs into the *s\_free* array. To free a block, check if *s\_nfree* is 50; if so, copy *s\_nfree* and the *s\_free* array into it, write it out, and set *s\_nfree* to 0. In any event set *s\_free*[*s\_nfree*] to the freed block's number and increment *s\_nfree*.

*S\_tfree* is the total free blocks available in the file system.

*S\_ninode* is the number of free i-numbers in the *s\_inode* array. To allocate an i-node: if *s\_ninode* is greater than 0, decrement it and return *s\_inode*[*s\_ninode*]. If it was 0, read the i-list and place the numbers of all free i-nodes (up to 100) into the *s\_inode* array, then try again. To free an i-node, provided *s\_ninode* is less than 100, place its number into *s\_inode*[*s\_ninode*] and increment *s\_ninode*. If *s\_ninode* is already 100, do not bother to enter the freed i-node into any table. This list of i-nodes is only to speed up the allocation process; the information as to whether the i-node is really free or not is maintained in the i-node itself.

*S\_tinode* is the total free i-nodes available in the file system.

*S\_flock* and *s\_iloc* are flags maintained in the core copy of the file system while it is mounted and their values on disk are immaterial. The value of *s\_fmmod* on disk is likewise immaterial; it is used as a flag to indicate that the super-block has changed and should be copied to the disk during the next periodic update of file system information.

*S\_ronly* is a read-only flag to indicate write-protection.

*S\_time* is the last time the super-block of the file system was changed, and is the number of seconds that have elapsed since 00:00 Jan. 1, 1970 (GMT). During a reboot, the *s\_time* of the super-block for the root file system is used to set the system's idea of the time.

*S\_fname* is the name of the file system and *s\_fpack* is the name of the pack.

I-numbers begin at 1, and the storage for i-nodes begins in block 2. Also, i-nodes are 64 bytes long. I-node 1 is reserved for future use. I-node 2 is reserved for the root directory of the file system, but no other i-number has a built-in meaning. Each i-node represents one file. For the format of an i-node and its flags, see *inode(4)*.

**FILES**

/usr/include/sys/filsys.h  
/usr/include/sys/stat.h

**SEE ALSO**

*inode(4)*;  
*fsck(1M)*, *fsdb(1M)* in the Runtime System manual.

**This page intentionally left blank.**

## NAME

fspec — format specification in text files

## DESCRIPTION

It is sometimes convenient to maintain text files on the UNIX system with non-standard tabs, (i.e., tabs which are not set at every eighth column). Such files must generally be converted to a standard format, frequently by replacing all tabs with the appropriate number of spaces, before they can be processed by UNIX system commands. A format specification occurring in the first line of a text file specifies how tabs are to be expanded in the remainder of the file.

A format specification consists of a sequence of parameters separated by blanks and surrounded by the brackets <: and :>. Each parameter consists of a keyletter, possibly followed immediately by a value. The following parameters are recognized:

**ttabs** The **t** parameter specifies the tab settings for the file. The value of **tabs** must be one of the following:

1. a list of column numbers separated by commas, indicating tabs set at the specified columns;
2. a — followed immediately by an integer *n*, indicating tabs at intervals of *n* columns;
3. a — followed by the name of a “canned” tab specification.

Standard tabs are specified by **t-8**, or equivalently, **t1,9,17,25**, etc. The canned tabs which are recognized are defined by the **tabs(1)** command.

**ssize** The **s** parameter specifies a maximum line size. The value of **size** must be an integer. Size checking is performed after tabs have been expanded, but before the margin is prepended.

**mmargin** The **m** parameter specifies a number of spaces to be prepended to each line. The value of **margin** must be an integer.

**d** The **d** parameter takes no value. Its presence indicates that the line containing the format specification is to be deleted from the converted file.

**e** The **e** parameter takes no value. Its presence indicates that the current format is to prevail only until another format specification is encountered in the file.

Default values, which are assumed for parameters not supplied, are **t-8** and **m0**. If the **s** parameter is not specified, no size checking is performed. If the first line of a file does not contain a format specification, the above defaults are assumed for the entire file. The following is an example of a line containing a format specification:

• <:t5,10,15 s72:> •

If a format specification can be disguised as a comment, it is not necessary to code the **d** parameter.

Several UNIX system commands correctly interpret the format specification for a file. Among them is *gath* [see *send(1C)*] which may be used to convert files to a standard format acceptable to other UNIX system commands.

## SEE ALSO

*ed(1)*, *newform(1)*, *send(1C)*, *tabs(1)* in the Runtime System manual.

# GETTYDEFS(4)

## NAME

gettydefs — speed and terminal settings used by *getty*

## DESCRIPTION

The */etc/gettydefs* file contains information used by *getty*(1M) to set up the speed and terminal settings for a line. It supplies information on what the *login* prompt should look like. It also supplies the speed to try next if the user indicates the current speed is not correct by typing a *<break>* character.

Each entry in */etc/gettydefs* has the following format:

label# initial-flags # final-flags # login-prompt #next-label

Each entry is followed by a blank line. The various fields can contain quoted characters of the form *\b*, *\n*, *\c*, etc., as well as *\nnn*, where *nnn* is the octal value of the desired character. The various fields are:

*label* This is the string against which *getty* tries to match its second argument. It is often the speed, such as **1200**, at which the terminal is supposed to run, but it need not be (see below).

*initial-flags* These flags are the initial *ioctl*(2) settings to which the terminal is to be set if a terminal type is not specified to *getty*. The flags that *getty* understands are the same as the ones listed in */usr/include/sys/termio.h* [see *termio*(7)]. Normally only the speed flag is required in the *initial-flags*. *Getty* automatically sets the terminal to raw input mode and takes care of most of the other flags. The *initial-flag* settings remain in effect until *getty* executes *login*(1).

*final-flags* These flags take the same values as the *initial-flags* and are set just prior to *getty* executes *login*. The speed flag is again required. The composite flag **SANE** takes care of most of the other flags that need to be set so that the processor and terminal are communicating in a rational fashion. The other two commonly specified *final-flags* are **TAB3**, so that tabs are sent to the terminal as spaces, and **HUPCL**, so that the line is hung up on the final close.

*login-prompt* This entire field is printed as the *login-prompt*. Unlike the above fields where white space is ignored (a space, tab or new-line), they are included in the *login-prompt* field.

*next-label* If this entry does not specify the desired speed, indicated by the user typing a *<break>* character, then *getty* will search for the entry with *next-label* as its *label* field and set up the terminal for those settings. Usually, a series of speeds are linked together in this fashion, into a closed set; For instance, **2400** linked to **1200**, which in turn is linked to **300**, which finally is linked to **2400**.

If *getty* is called without a second argument, then the first entry of */etc/gettydefs* is used, thus making the first entry of */etc/gettydefs* the default entry. It is also used if *getty* cannot find the specified *label*. If */etc/gettydefs* itself is missing, there is one entry built into the command which will bring up a terminal at **300** baud.

It is strongly recommended that after making or modifying */etc/gettydefs*, it be run through *getty* with the *check* option to be sure there are no errors.

## FILES

*/etc/gettydefs*

**SEE ALSO**

`ioctl(2)`.

`getty(1M)`, `login(1)`, `termio(7)` in the Runtime System manual.

## GPS(4)

### NAME

gps — graphical primitive string, format of graphical files

### DESCRIPTION

GPS is a format used to store graphical data. Several routines have been developed to edit and display GPS files on various devices. Also, higher level graphics programs such as *plot* [in *stat*(1G)] and *vloc* [in *toc*(1G)] produce GPS format output files.

A GPS is composed of five types of graphical data or primitives.

### GPS PRIMITIVES

- lines** The *lines* primitive has a variable number of points from which zero or more connected line segments are produced. The first point given produces a *move* to that location. (A *move* is a relocation of the graphic cursor without drawing.) Successive points produce line segments from the previous point. Parameters are available to set *color*, *weight*, and *style* (see below).
- arc** The *arc* primitive has a variable number of points to which a curve is fit. The first point produces a *move* to that point. If only two points are included, a line connecting the points will result; if three points a circular arc through the points is drawn; and if more than three, lines connect the points. (In the future, a spline will be fit to the points if they number greater than three.) Parameters are available to set *color*, *weight*, and *style*.
- text** The *text* primitive draws characters. It requires a single point which locates the center of the first character to be drawn. Parameters are *color*, *font*, *textsize*, and *textangle*.
- hardware** The *hardware* primitive draws hardware characters or gives control commands to a hardware device. A single point locates the beginning location of the *hardware* string.
- comment** A *comment* is an integer string that is included in a GPS file but causes nothing to be displayed. All GPS files begin with a comment of zero length.

### GPS PARAMETERS

- color** *Color* is an integer value set for *arc*, *lines*, and *text* primitives.
- weight** *Weight* is an integer value set for *arc* and *lines* primitives to indicate line thickness. The value **0** is narrow weight, **1** is bold, and **2** is medium weight.
- style** *Style* is an integer value set for *lines* and *arc* primitives to give one of the five different line styles that can be drawn on TEKTRONIX 4010 series storage tubes. They are:
- 0** solid
  - 1** dotted
  - 2** dot dashed
  - 3** dashed
  - 4** long dashed
- font** An integer value set for *text* primitives to designate the text font to be used in drawing a character string. (Currently *font* is expressed as a four-bit *weight* value followed by a four-bit *style* value.)
- textsize** *Textsize* is an integer value used in *text* primitives to express the size of the characters to be drawn. *Textsize* represents the height of characters in absolute *universe-units* and is stored at one-fifth this value in the size-orientation (*so*) word (see below).



**textangle** *Textangle* is a signed integer value used in *text* primitives to express rotation of the character string around the beginning point. *Textangle* is expressed in degrees from the positive x-axis and can be a positive or negative value. It is stored in the size-orientation (*so*) word as a value 256/360 of it's absolute value.

#### ORGANIZATION

GPS primitives are organized internally as follows:

<b>lines</b>	<i>cw points sw</i>
<b>arc</b>	<i>cw points sw</i>
<b>text</b>	<i>cw point sw so [string]</i>
<b>hardware</b>	<i>cw point [string]</i>
<b>comment</b>	<i>cw [string]</i>

**cw** *Cw* is the control word and begins all primitives. It consists of four bits that contain a primitive-type code and twelve bits that contain the word-count for that primitive.

**point(s)** *Point(s)* is one or more pairs of integer coordinates. *Text* and *hardware* primitives only require a single *point*. *Point(s)* are values within a Cartesian plane or *universe* having 64K (-32K to +32K) points on each axis.

**sw** *Sw* is the style-word and is used in *lines*, *arc*, and *text* primitives. For all three, eight bits contain *color* information. In *arc* and *lines* eight bits are divided as four bits *weight* and four bits *style*. In the *text* primitive eight bits of *sw* contain the *font*.

**so** *So* is the size-orientation word used in *text* primitives. Eight bits contain text size and eight bits contain text rotation.

**string** *String* is a null-terminated character string. If the string does not end on a word boundary, an additional null is added to the GPS file to insure word-boundary alignment.

#### SEE ALSO

*graphics(1G)*, *stat(1G)*, *toc(1G)* in the Runtime System manual.

## GROUP(4)

### NAME

group — group file

### DESCRIPTION

*Group* contains for each group the following information:

- group name
- encrypted password
- numerical group ID
- comma-separated list of all users allowed in the group

This is an ASCII file. The fields are separated by colons; each group is separated from the next by a new-line. If the password field is null, no password is demanded.

This file resides in directory */etc*. Because of the encrypted passwords, it can and does have general read permission and can be used, for example, to map numerical group ID's to names.

### FILES

*/etc/group*

### SEE ALSO

*crypt(3C)*, *passwd(4)*.  
*newgrp(1)*, *passwd(1)* in the Runtime System manual.

## NAME

inittab — script for the init process

## DESCRIPTION

The *inittab* file supplies the script to *init*'s role as a general process dispatcher. The process that constitutes the majority of *init*'s process dispatching activities is the line process */etc/getty* that initiates individual terminal lines. Other processes typically dispatched by *init* are daemons and the shell.

The *inittab* file is composed of entries that are position dependent and have the following format:

```
id:rstate:action:process
```

Each entry is delimited by a new-line; however, a backslash (\) preceding a new-line indicates a continuation of the entry. Up to 512 characters per entry are permitted. Comments may be inserted in the *process* field using the *sh*(1) convention for comments. Comments for lines that spawn *gettys* are displayed by the *who*(1) command. It is expected that they will contain some information about the line such as the location. There are no limits (other than maximum entry size) imposed on the number of entries within the *inittab* file. The entry fields are:

- id** This is one or two characters used to uniquely identify an entry.
- rstate** This defines the *run-level* in which this entry is to be processed. *Run-levels* effectively correspond to a configuration of processes in the system. That is, each process spawned by *init* is assigned a *run-level* or *run-levels* in which it is allowed to exist. The *run-levels* are represented by a number ranging from 0 through 6. As an example, if the system is in *run-level 1*, only those entries having a 1 in the *rstate* field will be processed. When *init* is requested to change *run-levels*, all processes which do not have an entry in the *rstate* field for the target *run-level* will be sent the warning signal (SIGTERM) and allowed a 20-second grace period before being forcibly terminated by a kill signal (SIGKILL). The *rstate* field can define multiple *run-levels* for a process by selecting more than one *run-level* in any combination from 0-6. If no *run-level* is specified, then the process is assumed to be valid at all *run-levels 0-6*. There are three other values, a, b and c, which can appear in the *rstate* field, even though they are not true *run-levels*. Entries which have these characters in the *rstate* field are processed only when the *telinit* [see *init*(1M)] process requests them to be run (regardless of the current *run-level* of the system). They differ from *run-levels* in that *init* can never enter *run-level a, b* or *c*. Also, a request for the execution of any of these processes does not change the current *run-level*. Furthermore, a process started by an a, b or c command is not killed when *init* changes levels. They are only killed if their line in */etc/inittab* is marked off in the *action* field, their line is deleted entirely from */etc/inittab*, or *init* goes into the *SINGLE USER* state.
- action** Key words in this field tell *init* how to treat the process specified in the *process* field. The actions recognized by *init* are as follows:
- respawn** If the process does not exist then start the process, do not wait for its termination (continue scanning the *inittab* file), and when it dies restart the process. If the process currently exists then do nothing and continue scanning the *inittab* file.
- wait** Upon *init*'s entering the *run-level* that matches the entry's *rstate*, start the process and wait for its

- termination. All subsequent reads of the *inittab* file while *init* is in the same *run-level* will cause *init* to ignore this entry.
- once** Upon *init*'s entering a *run-level* that matches the entry's *rstate*, start the process, do not wait for its termination. When it dies, do not restart the process. If upon entering a new *run-level*, where the process is still running from a previous *run-level* change, the program will not be restarted.
- boot** The entry is to be processed only at *init*'s boot-time read of the *inittab* file. *Init* is to start the process, not wait for its termination; and when it dies, not restart the process. In order for this instruction to be meaningful, the *rstate* should be the default or it must match *init*'s *run-level* at boot time. This action is useful for an initialization function following a hardware reboot of the system.
- bootwait** The entry is to be processed only at *init*'s boot-time read of the *inittab* file. *Init* is to start the process, wait for its termination and, when it dies, not restart the process.
- powerfail** Execute the process associated with this entry only when *init* receives a power fail signal [SIGPWR see *signal(2)*].
- powerwait** Execute the process associated with this entry only when *init* receives a power fail signal (SIGPWR) and wait until it terminates before continuing any processing of *inittab*.
- off** If the process associated with this entry is currently running, send the warning signal (SIGTERM) and wait 20 seconds before forcibly terminating the process via the kill signal (SIGKILL). If the process is nonexistent, ignore the entry.
- ondemand** This instruction is really a synonym for the **respawn** action. It is functionally identical to **respawn** but is given a different keyword in order to divorce its association with *run-levels*. This is used only with the **a**, **b** or **c** values described in the *rstate* field.
- initdefault** An entry with this *action* is only scanned when *init* initially invoked. *Init* uses this entry, if it exists, to determine which *run-level* to enter initially. It does this by taking the highest *run-level* specified in the *rstate* field and using that as its initial state. If the *rstate* field is empty, this is interpreted as **0123456** and so *init* will enter *run-level 6*. Also, the **initdefault** entry cannot specify that *init* start in the *SINGLE USER* state. Additionally, if *init* does not find an **initdefault** entry in */etc/inittab*, then it will request an initial *run-level* from the user at reboot time.
- sysinit** Entries of this type are executed before *init* tries to access the console. It is expected that this entry will be only used to initialize devices on which *init* might try to ask the *run-level* question. These entries are executed and waited for before continuing.

*process* This is a *sh* command to be executed. The entire **process** field is prefixed with *exec* and passed to a forked *sh* as *sh -c 'exec command'*. For this reason, any legal *sh* syntax can appear in the *process* field. Comments can be inserted with the ; *#comment* syntax.

**FILES**

/etc/inittab

**SEE ALSO**

*exec*(2), *open*(2), *signal*(2).

*getty*(1M), *init*(1M), *sh*(1), *who*(1) in the Runtime System manual.

## INODE(4)

### NAME

inode - format of an i-node

### SYNOPSIS

```
#include <sys/types.h>
#include <sys/ino.h>
```

### DESCRIPTION

An i-node for a plain file or directory in a file system has the following structure defined by `<sys/ino.h>`.

```
/* Inode structure as it appears on a disk block. */
struct dinode
{
    ushort di_mode; /* mode and type of file */
    short di_nlink; /* number of links to file */
    ushort di_uid; /* owner's user id */
    ushort di_gid; /* owner's group id */
    off_t di_size; /* number of bytes in file */
    char di_addr[40]; /* disk block addresses */
    time_t di_atime; /* time last accessed */
    time_t di_mtime; /* time last modified */
    time_t di_ctime; /* time of last file status change */
};
/*
 * the 40 address bytes:
 * 39 used; 13 addresses
 * of 3 bytes each.
 */
```

For the meaning of the defined types `off_t` and `time_t` see `types(5)`.

### FILES

`/usr/include/sys/ino.h`

### SEE ALSO

`stat(2)`, `fs(4)`, `types(5)`.

**NAME** issue - issue identification file

**DESCRIPTION**

The file `/etc/issue` contains the *issue* or project identification to be printed as a login prompt. This is an ASCII file which is read by program *getty* and then written to any terminal spawned or respawned from the *lines* file.

**FILES**

`/etc/issue`

**SEE ALSO**

`login(1)` in the Runtime System manual.

## LDFCN(4)

### NAME

ldfcn - common object file access routines

### SYNOPSIS

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

### DESCRIPTION

The common object file access routines are a collection of functions for reading an object file that is in common object file format. Although the calling program must know the detailed structure of the parts of the object file that it processes, the routines effectively insulate the calling program from knowledge of the overall structure of the object file.

The interface between the calling program and the object file access routines is based on the defined type **LDFILE**, defined as **struct ldfile**, declared in the header file **ldfcn.h**. The primary purpose of this structure is to provide uniform access to both simple object files and to object files that are members of an archive file.

The function *ldopen*(3X) allocates and initializes the **LDFILE** structure and returns a pointer to the structure to the calling program. The fields of the **LDFILE** structure may be accessed individually through macros defined in **ldfcn.h** and contain the following information:

**LDFILE**           \*ldptr;

**TYPE**(ldptr)    The file type, used to distinguish between archive members and simple object files. If the file is a simple object file, **TYPE**(ldptr) will contain the file magic number [see *filehdr*(4)]. If the file is an archive, **TYPE**(ldptr) will be **ARTYPE**, defined in *ldfcn.h*.

**IOPTR**(ldptr)   The file pointer returned by *fopen* and used by the standard input/output functions.

**OFFSET**(ldptr) The file address of the beginning of the object file; the offset is nonzero if the object file is a member of an archive file.

**HEADER**(ldptr)

The file header structure of the object file.

The object file access functions themselves may be divided into four categories:

- (1) functions that open or close an object file

*ldopen*(3X) and *ldaopen*  
    open a common object file  
*ldclose*(3X) and *ldaclose*  
    close a common object file

- (2) functions that read header or symbol table information

*ldahread*(3X)  
    read the archive header of a member of an archive file  
*ldfthread*(3X)  
    read the file header of a common object file  
*ldshread*(3X) and *ldnshread*  
    read a section header of a common object file  
*ldtbread*(3X)  
    read a symbol table entry of a common object file



*ldgetname(3X)*

retrieve a symbol name from a symbol table entry or from the string table

(3) functions that position an object file at (seek to) the start of the section, relocation, or line number information for a particular section.

*ldohseek(3X)*

seek to the optional file header of a common object file

*ldsseek(3X)* and *ldnsseek*

seek to a section of a common object file

*ldrseek(3X)* and *ldnrseek*

seek to the relocation information for a section of a common object file

*ldlseek(3X)* and *ldnlseek*

seek to the line number information for a section of a common object file

*ldtbseek(3X)*

seek to the symbol table of a common object file

(4) the function *ldtbindex(3X)* which returns the index of a particular common object file symbol table entry

These functions are described in detail in their respective manual pages.

All the functions except *ldopen*, *ldgetname*, *ldaopen* and *ldtbindex* return either **SUCCESS** or **FAILURE**, both constants defined in **ldfcn.h**. *Ldopen* and *ldaopen* both return pointers to a **LDFILE** structure.

**MACROS**

Additional access to an object file is provided through a set of macros defined in **ldfcn.h**. These macros parallel the standard input/output file reading and manipulating functions, translating a reference of the **LDFILE** structure into a reference to its file descriptor field.

The following macros are provided:

GETC(ldptr)  
 FGETC(ldptr)  
 GETW(ldptr)  
 UNGETC(c, ldptr)  
 FGETS(s, n, ldptr)  
 FREAD((char \*) ptr, sizeof (\*ptr), nitems, ldptr)  
 FSEEK(ldptr, offset, ptrname)  
 FTELL(ldptr)  
 REWIND(ldptr)  
 FEOF(ldptr)  
 FERROR(ldptr)  
 FILENO(ldptr)  
 SETBUF(ldptr, buf)  
 STROFFSET(ldptr)

The **STROFFSET** macro calculates the address of the string table in a UNIX System V Release 2.0 object file. See the manual entries for the corresponding standard input/output library functions for details on the use of the rest of the macros.

The program must be loaded with the object file access routine library **libld.a**.

**CAVEAT**

The macro **FSEEK** defined in the header file **ldfcn.h** translates into a

## LDFCN(4)

call to the standard input/output function *fseek(3S)*. **FSEEK** should not be used to seek from the end of an archive file since the end of an archive file may not be the same as the end of one of its object file members!

### SEE ALSO

*fseek(3S)*, *ldahread(3X)*, *ldclose(3X)*, *ldgetname(3X)*, *ldhread(3X)*,  
*ldhread(3X)*, *ldlseek(3X)*, *ldohseek(3X)*, *ldopen(3X)*, *ldrseek(3X)*,  
*ldlseek(3X)*, *ldshread(3X)*, *ldtbindex(3X)*, *ldtbread(3X)*, *ldtbseek(3X)*,  
*a.out(4)*, *ar(4)*, *filehdr(4)*,  
*Common Object File Format*.

**NAME**

linenum — line number entries in a common object file

**SYNOPSIS**

```
#include <linenum.h>
```

**DESCRIPTION**

Compilers based on *pcc* generate an entry in the object file for each C source line on which a breakpoint is possible [when invoked with the *-g* option; see *cc(1)*]. Users can then reference line numbers when using the appropriate software test system [see *sdb(1)*]. The structure of these line number entries appears below.

```
struct lineno
{
    union
    {
        long   l_symndx ;
        long   l_paddr ;
    }         l_addr ;
    unsigned short l_inno ;
};
```

Numbering starts with one for each function. The initial line number entry for a function has *l\_inno* equal to zero, and the symbol table index of the function's entry is in *l\_symndx*. Otherwise, *l\_inno* is non-zero, and *l\_paddr* is the physical address of the code for the referenced line. Thus the overall structure is the following:

<i>l_addr</i>	<i>l_inno</i>
function symtab index	0
physical address	line
physical address	line
...	
function symtab index	0
physical address	line
physical address	line
...	

**SEE ALSO**

*a.out(4)*,  
*cc(1)*, *sdb(1)* in the Runtime System manual.

S-4

# MASTER(4)

## NAME

master — master device information table

## DESCRIPTION

This file is used by the *config(1M)* program to obtain device information that enables it to generate the configuration file. *Master* contains lines of various forms for controlling the configuration of hardware devices, software drivers, parameters, and aliases.

In Part 1, hardware devices and software drivers are defined as follows:

- Field 1: device name (8 chars maximum).
- Field 2: interrupt vector
- Field 3: functions for this device:
  - o** open handler
  - c** close handler
  - r** read handler
  - w** write handler
  - i** ioctl handler
  - s** startup routine
  - f** fork
  - e** exec
  - x** exit
- Field 4: element characteristics:
  - o** specify only once
  - s** suppress count field
  - r** required device
  - b** block device
  - c** character device
  - t** device is a tty
- Field 5: handler prefix
- Field 6: major device number if block-type device
- Field 7: major device number if character-type device
- Field 8: number of sub-devices per device
- Field 9: configuration table structure

Part 2 contains lines with 2 fields each:

- Field 1: alias name of device (8 chars. maximum).
- Field 2: reference name of device (8 chars. maximum; specified in part 1).

Part 3 contains lines with 2 or 3 fields each:

- Field 1: parameter name (as it appears in description file; 20 chars. maximum)
- Field 2: parameter name (as it appears in the *conf.c* file; 20 chars. maximum)
- Field 3: default parameter value (20 chars. maximum; parameter specification is required if this field is omitted)

## SEE ALSO

*config(1M)*, *sysdef(1M)* in the Runtime System manual.

**NAME**

mnttab — mounted file system table

**SYNOPSIS**

```
#include <mnttab.h>
```

**DESCRIPTION**

*Mnttab* resides in directory */etc* and contains a table of devices, mounted by the *mount(1M)* command, in the following structure as defined by *<mnttab.h>*:

```
struct mnttab {
    char    mt_dev[32];
    char    mt_filsys[32];
    short   mt_ro_flg;
    time_t  mt_time;
};
```

Each entry is 70 bytes in length; the first 32 bytes are the null-padded name of the place where the *special file* is mounted; the next 32 bytes represent the null-padded root name of the mounted special file; the remaining 6 bytes contain the mounted *special file*'s read/write permissions and the date on which it was mounted.

The maximum number of entries in *mnttab* is based on the system parameter *NMOUNT* located in */usr/src/uts/cf/conf.c*, which defines the number of allowable mounted special files.

**SEE ALSO**

*mount(1M)*, *setmnt(1M)* in the Runtime System manual.

## PASSWD(4)

### NAME

passwd — password file

### DESCRIPTION

*Passwd* contains for each user the following information:

- login name
- encrypted password
- numerical user ID
- numerical group ID
- GCOS job number, box number, optional GCOS user ID
- initial working directory
- program to use as Shell

This is an ASCII file. Each field within each user's entry is separated from the next by a colon. The GCOS field is used only when communicating with that system, and in other installations can contain any desired information. Each user is separated from the next by a new-line. If the password field is null, no password is demanded; if the Shell field is null, the Shell itself is used.

This file resides in directory */etc*. Because of the encrypted passwords, it can and does have general read permission and can be used, for example, to map numerical user IDs to names.

The encrypted password consists of 13 characters chosen from a 64-character alphabet (*, /, 0-9, A-Z, a-z*), except when the password is null, in which case the encrypted password is also null. Password aging is effected for a particular user if this encrypted password in the password file is followed by a comma and a non-null string of characters from the above alphabet. (Such a string must be introduced in the first instance by the super-user.)

The first character of the age, *M* say, denotes the maximum number of weeks for which a password is valid. Users who attempt to log in after their passwords have expired will be forced to supply a new one. The next character, *m* say, denotes the minimum period in weeks which must expire before the password may be changed. The remaining characters define the week (counted from the beginning of 1970) when the password was last changed. (A null string is equivalent to zero.) *M* and *m* have numerical values in the range 0-63 that correspond to the 64-character alphabet shown above (i.e., */* = 1 week; *z* = 63 weeks). If *m* = *M* = 0 (derived from the string *.or.*) users will be forced to change their passwords the next time they log in (and the "age" will disappear from their entries in the password file). If *m* > *M* (signified, e.g., by the string *.N*) only the super-user will be able to change the password.

### FILES

*/etc/passwd*

### SEE ALSO

*a64l(3C)*, *crypt(3C)*, *getpwent(3C)*, *group(4)*,  
*login(1)*, *passwd(1)* in the Runtime System manual.

## NAME

plot — graphics interface

## DESCRIPTION

Files of this format are produced by routines described in *plot(3X)* and are interpreted for various devices by commands described in *tplot(1G)*. A graphics file is a stream of plotting instructions. Each instruction consists of an ASCII letter usually followed by bytes of binary information. The instructions are executed in order. A point is designated by four bytes representing the *x* and *y* values; each value is a signed integer. The last designated point in an *l*, *m*, *n*, or *p* instruction becomes the “current point” for the next instruction.

Each of the following descriptions begins with the name of the corresponding routine in *plot(3X)*.

- m** move: The next four bytes give a new current point.
- n** cont: Draw a line from the current point to the point given by the next four bytes. See *tplot(1G)*.
- p** point: Plot the point given by the next four bytes.
- l** line: Draw a line from the point given by the next four bytes to the point given by the following four bytes.
- t** label: Place the following ASCII string so that its first character falls on the current point. The string is terminated by a new-line.
- e** erase: Start another frame of output.
- f** linemod: Take the following string, up to a new-line, as the style for drawing further lines. The styles are “dotted”, “solid”, “longdashed”, “short-dashed”, and “dottedashed”. Effective only for the **-T4014** and **-Tvers** options of *tplot(1G)* (TEKTRONIX 4014 terminal and Versatec plotter).
- s** space: The next four bytes give the lower left corner of the plotting area; the following four give the upper right corner. The plot will be magnified or reduced to fit the device as closely as possible.

Space settings that exactly fill the plotting area with unity scaling appear below for devices supported by the filters of *tplot(1G)*. The upper limit is just outside the plotting area. In every case the plotting area is taken to be square; points outside may be displayable on devices whose face is not square.

DASI 300	space(0, 0, 4096, 4096);
DASI 300s	space(0, 0, 4096, 4096);
DASI 450	space(0, 0, 4096, 4096);
TEKTRONIX 4014	space(0, 0, 3120, 3120);
Versatec plotter	space(0, 0, 2048, 2048);

## SEE ALSO

*plot(3X)*, *gps(4)*, *term(5)*,  
*graph(1G)*, *tplot(1G)* in the Runtime System manual.

## WARNING

The plotting library *plot(3X)* and the curses library *curses(3X)* both use the names *erase()* and *move()*. The curses versions are macros. If you need both libraries, put the *plot(3X)* code in a different source file than the *curses(3X)* code, and/or `#undef move()` and `erase()` in the *plot(3X)* code.

## **PNCH(4)**

### **NAME**

pnch — file format for card images

### **DESCRIPTION**

The PNCH format is a convenient representation for files consisting of card images in an arbitrary code.

A PNCH file is a simple concatenation of card records. A card record consists of a single control byte followed by a variable number of data bytes. The control byte specifies the number (which must lie in the range 0-80) of data bytes that follow. The data bytes are 8-bit codes that constitute the card image. If there are fewer than 80 data bytes, it is understood that the remainder of the card image consists of trailing blanks.

### **SEE ALSO**

send(1C) in the Runtime System manual.



## NAME

profile - setting up an environment at login time

## DESCRIPTION

If your login directory contains a file named **.profile**, that file will be executed (via **exec .profile**) before your session begins; **.profiles** are handy for setting exported environment variables and terminal modes. If the file **/etc/profile** exists, it will be executed for every user before the **.profile**. The following example is typical (except for the comments):

```
# Make some environment variables global
export MAIL PATH TERM
# Set file creation mask
umask 22
# Tell me when new mail comes in
MAIL=/usr/mail/myname
# Add my /bin directory to the shell search sequence
PATH=$PATH:$HOME/bin
# Set terminal type
echo "terminal: \c"
read TERM
case $TERM in
    300)          stty cr2 nl0 tabs; tabs;;
    300s)         stty cr2 nl0 tabs; tabs;;
    450)          stty cr2 nl0 tabs; tabs;;
    hp)           stty cr0 nl0 tabs; tabs;;
    745 | 735)    stty cr1 nl1 -tabs; TERM=745;;
    43)           stty cr1 nl0 -tabs;;
    4014 | tek)   stty cr0 nl0 -tabs ff1; TERM=4014; echo "\33;" ;;
    *)           echo "$TERM unknown" ;;
esac
```

## FILES

\$HOME/.profile  
/etc/profile

## SEE ALSO

environ(5), term(5),  
env(1), login(1), mail(1), sh(1), stty(1), su(1) in the Runtime System manual.

## RELOC(4)

### NAME

reloc - relocation information for an object file

### SYNOPSIS

```
#include <reloc.h>
```

### DESCRIPTION

Object files have one relocation entry for each relocatable reference in the text or data. If relocation information is present, it will be in the following format.

```
struct reloc
{
    long    r_vaddr;    /* (virtual) address of */
                    /* reference */
    long    r_symndx;   /* index into symbol table */
    short   r_type;     /* relocation type */
};

#define R_ABS      0
#define R_DIR16   01
#define R_REL16   02
#define R_IND16   03
#define R_OFF8    07
#define R_OFF16   10
#define R_SEG12   11
#define R_AUX     13
```

As the link editor reads each input section and performs relocation, the relocation entries are read. They direct how references found within the input section are treated.

- R\_ABS**      The reference is absolute, and no relocation is necessary. The entry will be ignored.
- R\_DIR16**    A direct, 16-bit reference to a symbol's virtual address.
- R\_REL16**    A "PC-relative", 16-bit reference to a symbol's virtual address. Relative references occur in instructions such as jumps and calls. The actual address used is obtained by adding a constant to the value of the program counter at the time the instruction is executed.
- R\_IND16**    An indirect, 16-bit reference through a transfer vector. The instruction contains the virtual address of the transfer vector, where the actual address of the referenced word is stored.
- R\_OFF8**     A direct, 16-bit reference to the low-order 8 bits of a 20-bit virtual address. The 16-bit field has its high-order 8 bits forced to zero.
- R\_OFF16**    A direct, 16-bit reference to the low-order 16 bits of a 32-bit virtual address. The 16-bit field is treated as an unsigned integer. This relocation type is used when a (16-bit) constant modifies the virtual address.
- R\_SEG12**    A direct, 16-bit reference to the high-order 16 bits of a 32-bit virtual address.
- R\_AUX**      An "auxiliary entry", generated to permit the correct processing of relocation entries of type R\_SEG12. Each R\_SEG12 entry is followed immediately by a R\_AUX entry.

Other relocation types will be defined as they are needed.

Relocation entries are generated automatically by the assembler and automatically utilized by the link editor. A link editor option exists for removing the relocation entries from an object file.

**SEE ALSO**

a.out(4), syms(4).

ld(1), strip(1) in the Runtime System manual.

**This page intentionally left blank.**

**NAME**

scsfile — format of SCCS file

**DESCRIPTION**

An SCCS file is an ASCII file. It consists of six logical parts: the *checksum*, the *delta table* (contains information about each delta), *user names* (contains login names and/or numerical group IDs of users who may add deltas), *flags* (contains definitions of internal keywords), *comments* (contains arbitrary descriptive information about the file), and the *body* (contains the actual text lines intermixed with control lines).

Throughout an SCCS file there are lines which begin with the ASCII SOH (start of heading) character (octal 001). This character is hereafter referred to as *the control character* and will be represented graphically as @. Any line described below which is not depicted as beginning with the control character is prevented from beginning with the control character.

Entries of the form **DDDDD** represent a five-digit string (a number between 00000 and 99999).

Each logical part of an SCCS file is described in detail below.

*Checksum*

The checksum is the first line of an SCCS file. The form of the line is:  
**@bDDDDD**

The value of the checksum is the sum of all characters, except those of the first line. The **@b** provides a *magic number* of (octal) 064001.

*Delta table*

The delta table consists of a variable number of entries of the form:

```

@s DDDDD/DDDDD/DDDDD
@d <type> <SCCS ID> yr/ma/da hr:mi:se <pgmr> DDDDD DDDDD
@i DDDDD ...
@x DDDDD ...
@g DDDDD ...
@m <MR number>
.
.
.
@c <comments> ...
.
.
.
@e

```

The first line (**@s**) contains the number of lines inserted/deleted/unchanged, respectively. The second line (**@d**) contains the type of the delta (currently, normal: **D**, and removed: **R**), the SCCS ID of the delta, the date and time of creation of the delta, the login name corresponding to the real user ID at the time the delta was created, and the serial numbers of the delta and its predecessor, respectively.

The **@i**, **@x**, and **@g** lines contain the serial numbers of deltas included, excluded, and ignored, respectively. These lines are optional.

The @m lines (optional) each contain one MR number associated with the delta; the @c lines contain comments associated with the delta.

The @e line ends the delta table entry.

#### User names

The list of login names and/or numerical group IDs of users who may add deltas to the file, separated by new-lines. The lines containing these login names and/or numerical group IDs are surrounded by the bracketing lines @u and @U. An empty list allows anyone to make a delta. Any line starting with a ! prohibits the succeeding group or user from making deltas.

#### Flags

Keywords used internally [see *admin*(1) for more information on their use]. Each flag line takes the form:

```
@f <flag>    <optional text>
```

The following flags are defined:

```
@f t    <type of program>
@f v    <program name>
@f i    <keyword string>
@f b
@f m    <module name>
@f f    <floor>
@f c    <ceiling>
@f d    <default-sid>
@f n
@f j
@f l    <lock-releases>
@f q    <user defined>
@f z    <reserved for use in interfaces>
```

The **t** flag defines the replacement for the %Y% identification keyword. The **v** flag controls prompting for MR numbers in addition to comments; if the optional text is present it defines an MR number validity checking program. The **i** flag controls the warning/error aspect of the "No id keywords" message. When the **i** flag is not present, this message is only a warning; when the **i** flag is present, this message will cause a "fatal" error (the file will not be gotten, or the delta will not be made). When the **b** flag is present the **-b** keyletter may be used on the *get* command to cause a branch in the delta tree. The **m** flag defines the first choice for the replacement text of the %M% identification keyword. The **f** flag defines the "floor" release; the release below which no deltas may be added. The **c** flag defines the "ceiling" release; the release above which no deltas may be added. The **d** flag defines the default SID to be used when none is specified on a *get* command. The **n** flag causes *delta* to insert a "null" delta (a delta that applies *no* changes) in those releases that are skipped when a delta is made in a *new* release (e.g., when delta 5.1 is made after delta 2.7, releases 3 and 4 are skipped). The absence of the **n** flag causes skipped releases to be completely empty. The **j** flag causes *get* to allow concurrent edits of the same base SID. The **l** flag defines a *list* of releases that are *locked* against editing [*get*(1) with the **-e** keyletter]. The **q** flag defines the replacement for the %Q% identification keyword. The **z** flag is used in certain specialized interface programs.

*Comments*

Arbitrary text is surrounded by the bracketing lines @t and @T. The comments section typically will contain a description of the file's purpose.

*Body*

The body consists of text lines and control lines. Text lines do not begin with the control character, control lines do. There are three kinds of control lines: *insert*, *delete*, and *end*, represented by:

**@I DDDDD**

**@D DDDDD**

**@E DDDDD**

respectively. The digit string is the serial number corresponding to the delta for the control line.

**SEE ALSO**

admin(1), delta(1), get(1), prs(1) in the Runtime System manual.  
 "Source Code Control System User Guide" in the Software Development System manual, Vol. I.

## SCNHDR(4)

### NAME

scnhdr - section header for a common object file

### SYNOPSIS

```
#include <scnhdr.h>
```

### DESCRIPTION

Every common object file has a table of section headers to specify the layout of the data within the file. Each section within an object file has its own header. The C structure appears below.

```
struct scnhdr
{
    char        s_name[SYMNMLEN]; /* section name */
    long        s_paddr; /* physical address */
    long        s_vaddr; /* virtual address */
    long        s_size; /* section size */
    long        s_scnptr; /* file ptr to raw data */
    long        s_relptr; /* file ptr to relocation */
    long        s_lnnoptr; /* file ptr to line numbers */
    unsigned short s_nreloc; /* # reloc entries */
    unsigned short s_nlnno; /* # line number entries */
    long        s_flags; /* flags */
};
```

File pointers are byte offsets into the file; they can be used as the offset in a call to *fseek*(3S). If a section is initialized, the file contains the actual bytes. An uninitialized section is somewhat different. It has a size, symbols defined in it, and symbols that refer to it. But it can have no relocation entries, line numbers, or data. Consequently, an uninitialized section has no raw data in the object file, and the values for *s\_scnptr*, *s\_relptr*, *s\_lnnoptr*, *s\_nreloc*, and *s\_nlnno* are zero.

### SEE ALSO

*fseek*(3S), *a.out*(4).  
*ld*(1) in the Runtime System manual.



Flag bits currently defined are:

```
#define FSA_ITER      0x0001 /* true if seg contains iter data */
#define FSA_HUGE     0x0002 /* true if huge data item in seg */
#define FSA_BSS      0x0004 /* true if seg contains implicit BSS */
#define FSA_SHARE    0x0008 /* true if seg is sharable */
#define FSA_EXPDOWN  0x0010 /* true if seg is expanddown */
#define FSA_SEG      0x8000 /* true if segment, else section */
```

### COFF Section Header

The following C struct declaration shows the format of the COFF section header:

```
struct scnhdr
{
    char          s_name[SYMNMLEN]; /* section name */
    long          s_paddr; /* physical address */
    long          s_vaddr; /* virtual address */
    long          s_size; /* section size */
    long          s_scnptr; /* file ptr to raw data */
    long          s_relptr; /* file ptr to relocation */
    long          s_innoptr; /* file ptr to line numbers */
    unsigned short s_nreloc; /* # reloc entries */
    unsigned short s_nlnno; /* # line number entries */
    long          s_flags; /* flags */
};
```

### Map of COFF to STL

The following table shows the mapping of the (virtual format) COFF section-header fields to their actual STL locations. The following conventions are used:

The fileheader is referred to as "filehdr" in COFF and "stlhdr" in STL.

The system header is referred to as "aouthdr" in COFF and "exthdr" (extended header) in STL.

The section headers are referred to as "scnhdr" in COFF and "sechdr" in STL.

The mapping of some fields varies depending on the type of segment, the MODEL of the file, and if the file is executable. Type is derived from sechdr->se\_type and sechdr->se\_flags, and MODEL from stlhdr->fh\_flags (see filebdr.h).

For text:

```
scnhdr->s_name      = ".text";
scnhdr->s_size      = sechdr->se_psize;
if (filehdr->f_flags & F_EXEC)
    scnhdr->s_paddr = (sechdr->se_num << 19) |
                    ((exthdr->e_csselector << 16) & 0x7000)
else
    scnhdr->s_paddr = 0;
scnhdr->s_vaddr     = 0;
scnhdr->s_scnptr    = sechdr->se_scnptr;
scnhdr->s_flags     = STYP_TEXT; /* 0x020 */
```

For data:

```
scnhdr->s_name      = ".data";
scnhdr->s_size      = sechdr->se_psize;
if (filehdr->f_flags & F_EXEC)
```

## SCNHDR(4)

```
        scnhdr->s_paddr = (sechdr->se_num << 19) |
            ((exthdr->e_csselector << 16) & 0x7000)
    else
        scnhdr->s_paddr = sechdr->se_psize(TEXT)
    if (filehdr->f_flags & F_EXEC)
        if (MODEL == M_SMALL)
            scnhdr->s_vaddr = exthdr->e_stksiz
        else
            scnhdr->s_vaddr = 0;
    else
        scnhdr->s_vaddr = scnhdr->s_size(TEXT)
    scnhdr->s_scnptr = sechdr->se_scnptr;
    scnhdr->s_flags = STYP_DATA; /* 0x040 */

For bss:
    scnhdr->s_name = ".bss";
    if ((MODEL == M_SMALL) && (filehdr->f_flags & FH_EXEC))
        scnhdr->s_size = sthdr->fh_nsisiz
    else
        scnhdr->s_size = sechdr->se_vsize
    if (filehdr->f_flags & F_EXEC)
        if (MODEL == M_SMALL)
            scnhdr->s_paddr = scnhdr->s_paddr(DATA)
                + scnhdr->s_size(DATA) + exthdr->e_stksiz
        else
            scnhdr->s_paddr = (sechdr->s_num << 19) |
                ((exthdr->e_csselector << 16) & 0x7000)
    else
        scnhdr->s_paddr = scnhdr->s_psize(TEXT) +
            scnhdr->s_psize(DATA)
    if (filehdr->f_flags & F_EXEC)
        if (MODEL == M_SMALL)
            scnhdr->s_vaddr = exthdr->e_stksiz +
                scnhdr->s_size(DATA) + scnhdr->s_paddr(DATA)
        else
            scnhdr->s_vaddr = 0;
    else
        scnhdr->s_vaddr = scnhdr->s_size(DATA) +
            scnhdr->s_vaddr(DATA)
    scnhdr->s_scnptr = 0;
    scnhdr->s_flags = STYP_BSS; /* 0x080 */
```

The rest of the fields in the section header are mapped the same way regardless of the type of section:

```
    if (filehdr->f_flags & FH_EXEC)
        s_relptr = 0
    else
        s_relptr = sechdr->se_relptr;
    scnhdr->s_innoptr = sechdr->se_innoptr;
    scnhdr->s_nreloc = sechdr->se_nreloc;
    scnhdr->s_nlnno = sechdr->se_nlnno;
```

### SEE ALSO

seek(3S), a.out(4),  
c286ld(1) in the Runtime System manual.

## NAME

syms - common object file symbol table format

## SYNOPSIS

```
#include <syms.h>
```

## DESCRIPTION

Common object files contain information to support *symbolic* software testing. Line number entries, *linenum*(4), and extensive symbolic information permit testing at the C source level. Every object file's symbol table is organized as shown below.

```
File name 1.
  Function 1.
    Local symbols for function 1.
  Function 2
    Local symbols for function 2.
  ...
  Static externs for file 1.
```

```
File name 2.
  Function 1.
    Local symbols for function 1.
  Function 2.
    Local symbols for function 2.
  ...
  Static externs for file 2.
```

...

## Global symbols.

The entry for a symbol is a fixed-length structure. The members of the structure hold the name (null padded), its value, and other information. The C structure is given below.

```
#define SYMNMLEN  8
#define FILNMLEN 14

struct syment
{
  union
  {
    char      _n_name[SYMNMLEN]; /* symbol name */
    struct
    {
      long    _n_zeroes; /* == 0L when in string table */
      long    _n_offset; /* location of name in table */
    } _n_n;
    char      *_n_nptr[2]; /* allows overlaying */
  } _n;
  long        n_value; /* value of symbol */
  short       n_snum; /* section number */
  unsigned short n_type; /* type and derived type */
  char        n_sclass; /* storage class */
  char        n_numaux; /* number of aux entries */
};

#define n_name      _n_n_name
#define n_zeroes    _n_n_n_zeroes
```

## SYMS(4)

```
#define n_offset  _n_n_n_n_offset
#define n_nptr   _n_n_nptr[1]
```

Meaningful values and explanations for them are given in both **syms.h** and *Common Object File Format Chapter*. Anyone who needs to interpret the entries should seek more information in these sources. Some symbols require more information than a single entry; they are followed by *auxiliary entries* that are the same size as a symbol entry. The format follows.

```
union auxent
{
    struct
    {
        long          x_tagndx;
        union
        {
            struct
            {
                unsigned short x_lno;
                unsigned short x_size;
            } x_lnsz;
            long          x_fsize;
        } x_misc;
        union
        {
            struct
            {
                long  x_lnoptr;
                long  x_endndx;
            } struct
            {
                unsigned short x_dimen[DIMNUM];
                x_ary;
                x_fcary;
            }
            unsigned short x_tvndx;
        }
        x_sym;
    }
    struct
    {
        char  x_fname[FILNMLEN];
    }
    x_file;
}
struct
{
    long  x_scnlen;
    unsigned short x_nreloc;
    unsigned short x_nlinno;
}
x_scn;

struct
{
    long          x_tvfill;
    unsigned short x_tvlen;
    unsigned short x_tvran[2];
}
x_tv;
};
```

Indexes of symbol table entries begin at *zero*.

## SEE ALSO

a.out(4), linenum(4),  
*Common Object File Format.*

## CAVEATS

On machines in which longs are equivalent to ints (3B20S computer, VAX), they are converted to ints in the compiler to minimize the complexity of the compiler code generator. Thus the information about which symbols are declared as longs and which as ints does not show up in the symbol table.

## TERM(4)

### NAME

term - format of compiled term file.

### SYNOPSIS

**term**

### DESCRIPTION

Compiled terminfo descriptions are placed under the directory **/usr/lib/terminfo**. In order to avoid a linear search of a huge UNIX system directory, a two-level scheme is used: **/usr/lib/terminfo/c/name** where *name* is the name of the terminal, and *c* is the first character of *name*. Thus, *act4* can be found in the file **/usr/lib/terminfo/a/act4**. Synonyms for the same terminal are implemented by multiple links to the same compiled file.

The format has been chosen so that it will be the same on all hardware. An 8-or-more-bits byte is assumed, but no assumptions about byte-ordering or sign extension are made.

The compiled file is created with the *compile* program, and read by the routine *setupterm*. Both of these pieces of software are part of *curses*(3X). The file is divided into six parts: the header, terminal names, Boolean flags, numbers, strings, and string table.

The header section begins the file. This section contains six short integers in the format described below. These integers are (1) the magic number (octal 0432); (2) the size, in bytes, of the names section; (3) the number of bytes in the Boolean section; (4) the number of short integers in the numbers section; (5) the number of offsets (short integers) in the strings section; (6) the size, in bytes, of the string table.

Short integers are stored in two 8-bit bytes. The first byte contains the least significant 8 bits of the value, and the second byte contains the most significant 8 bits. (Thus, the value represented is  $256 * \text{second} + \text{first}$ .) The value -1 is represented by 0377, 0377, other negative values are illegal. The -1 generally means that a capability is missing from this terminal. Note that this format corresponds to the hardware of the VAX and PDP-11. Machines where this does not correspond to the hardware read the integers as two bytes and compute the result.

The terminal names section comes next. It contains the first line of the terminfo description, listing the various names for the terminal, separated by the |" character. The section is terminated with an ASCII NUL character.

The Boolean flags have one byte for each flag. This byte is either 0 or 1 as the flag is present or absent. The capabilities are in the same order as the file <term.h>.

Between the Boolean section and the number section, a null byte will be inserted, if necessary, to ensure that the number section begins on an even byte. All short integers are aligned on a short word boundary.

The numbers section is similar to the flags section. Each capability takes up two bytes, and is stored as a short integer. If the value represented is -1, the capability is taken to be missing.

The strings section is also similar. Each capability is stored as a short integer, in the format above. A value of -1 means the capability is missing. Otherwise, the value is taken as an offset from the beginning of the string table. Special characters in  $\backslash X$  or  $\backslash c$  notation are stored in their interpreted form, not the printing representation. Padding information  $\$<nn>$  and parameter information  $\%x$  are stored intact in uninterpreted form.

The final section is the string table. It contains all the values of string capabilities referenced in the string section. Each string is null terminated.

Note that it is possible for *setupterm* to expect a different set of capabilities than are actually present in the file. Either the database may have been updated since *setupterm* has been recompiled (resulting in extra unrecognized entries in the file) or the program may have been recompiled more recently than the database was updated (resulting in missing entries). The routine *setupterm* must be prepared for both possibilities — this is why the numbers and sizes are included. Also, new capabilities must always be added at the end of the lists of Boolean, number, and string capabilities.

As an example, an octal dump of the description for the Microterm ACT 4 is included:

```
microterm|act4|microterm act iv,
  cr=^M, cudl=^J, ind=^J, bel=^G, am, cubl=^H,
  ed=^_, el="--, clear=^L, cup=^T%p1%c%p2%c,
  cols#80, lines#24, cuf1=^X, cuu1=^Z, home=^],

000 032 001      \0 025 \0 \b \0 212 \0      " \0 m i c r
020 o t e r m | a c t 4 ; m i c r o
040 t e r m a c t i v \0 \0 001 \0 \0
060 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0 \0
100 \0 \0 p \0 377 377 030 \0 377 377 377 377 377 377 377 377
120 377 377 377 377 \0 \0 002 \0 377 377 377 377 004 \0 006 \0
140 \b \0 377 377 377 377 \n \0 026 \0 030 \0 377 377 032 \0
160 377 377 377 377 034 \0 377 377 036 \0 377 377 377 377 377 377
200 377 377 377 377 377 377 377 377 377 377 377 377 377 377
*
520 377 377 377 377      \0 377 377 377 377 377 377 377 377 377
540 377 377 377 377 377 377 007 \0 \r \0 \f \0 036 \0 037 \0
560 024 % p 1 % c % p 2 % c \0 \n \0 035 \0
600 \b \0 030 \0 032 \0 \n \0
```

Some limitations: total compiled entries cannot exceed 4096 bytes. The name field cannot exceed 128 bytes.

#### FILES

/usr/lib/terminfo/\*/\* compiled terminal capability data base

#### SEE ALSO

curses(3X), terminfo(4).

## TERMINFO(4)

### NAME

terminfo — terminal capability data base

### SYNOPSIS

/usr/lib/terminfo/\*/\*

### DESCRIPTION

*Terminfo* is a data base describing terminals, used, e.g., by *vi*(1) and  *curses*(3X). Terminals are described in *terminfo* by giving a set of capabilities which they have, and by describing how operations are performed. Padding requirements and initialization sequences are included in *terminfo*.

Entries in *terminfo* consist of a number of “,” separated fields. White space after each “,” is ignored. The first entry for each terminal gives the names which are known for the terminal, separated by “|” characters. The first name given is the most common abbreviation for the terminal, the last name given should be a long name fully identifying the terminal, and all others are understood as synonyms for the terminal name. All names but the last should be in lowercase and contain no blanks; the last name may well contain uppercase and blanks for readability.

Terminal names (except for the last, verbose entry) should be chosen using the following conventions. The particular piece of hardware making up the terminal should have a root name chosen, thus “hp2621”. This name should not contain hyphens, except that synonyms may be chosen that do not conflict with other names. Modes that the hardware can be in, or user preferences, should be indicated by appending a hyphen and an indicator of the mode. Thus, a vt100 in 132 column mode would be vt100-w. The following suffixes should be used where possible:

Suffix	Meaning	Example
-w	Wide mode (more than 80 columns)	vt100-w
-am	With auto. margins (usually default)	vt100-am
-nam	Without automatic margins	vt100-nam
-n	Number of lines on the screen	aaa-60
-na	No arrow keys (leave them in local)	c100-na
-np	Number of pages of memory	c100-4p
-rv	Reverse video	c100-rv

### CAPABILITIES

The variable is the name by which the programmer (at the terminfo level) accesses the capability. The capname is the short name used in the text of the data base, and is used by a person updating the data base. The i.code is the two letter internal code used in the compiled data base, and always corresponds to the old **termcap** capability name.

Capability names have no hard length limit, but an informal limit of 5 characters has been adopted to keep them short and to allow the tabs in the source file caps to line up nicely. Whenever possible, names are chosen to be the same as or similar to the ANSI X3.64-1979 standard. Semantics are also intended to match those of the specification.

- (P) indicates that padding may be specified
- (G) indicates that the string is passed through tparm withparms as given (#i).
- (\*) indicates that padding may be based on the number of lines affected
- (#.) indicates the  $i^{\text{th}}$  parameter.



Variable	Cap-name	I. Code	Description
<b>Booleans</b>			
auto_left_margin,	bw	bw	cb1 wraps from column 0 to last column
auto_right_margin,	am	am	Terminal has automatic margins
beehive_glitch,	xsb	xb	Beehive (f1=escape, f2=ctrl C)
ceol_standout_glitch,	xhp	xs	Standout not erased by overwriting (hp)
eat_newline_glitch,	xenl	xn	new-line ignored after 80 cols (Concept)
erase_overstrike,	eo	eo	Can erase overstrikes with a blank
generic_type,	gn	gn	Generic line type (e.g., dialup, switch).
hard_copy,	hc	hc	Hardcopy terminal
has_meta_key,	km	km	Has a meta key (shift, sets parity bit)
has_status_line,	hs	hs	Has extra "status line"
insert_null_glitch,	in	in	Insert mode distinguishes nulls
memory_above,	da	da	Display may be retained above the screen
memory_below,	db	db	Display may be retained below the screen
move_insert_mode,	mir	mi	Safe to move while in insert mode
move_standout_mode,	msgr	ms	Safe to move in standout modes
over_strike,	os	os	Terminal overstrikes
status_line_esc_ok,	eslok	es	Escape can be used on the status line
telcray_glitch,	xt	xt	Tabs ruin, magic so char (Telcray 1061)
tilde_glitch,	hz	hz	Hazeltine; can not print '~s
transparent_underline,	ul	ul	underline character overstrikes
xon_xoff,	xon	xo	Terminal uses xon/xoff handshaking
<b>Numbers:</b>			
columns,	cols	co	Number of columns in a line
init_tabs,	it	it	Tabs initially every # spaces
lines,	lines	lj	Number of lines on screen or page
lines_of_memory,	lm	lm	Lines of memory if > lines. 0 means varies
magic_cookie_glitch,	xmc	sg	Number of blank chars left by smso or rmso
padding_baud_rate,	pb	pb	Lowest baud where cr/nl padding is needed
virtual_terminal,	vt	vt	Virtual terminal number (UNIX system)
width_status_line,	wsl	ws	No. columns in status line
<b>Strings:</b>			
back_tab,	cbt	bt	Back tab (P)
bell,	bel	bl	Audible signal (bell) (P)
carriage_return,	cr	cr	Carriage return (P*)
change_scroll_region,	csr	cs	change to lines #1 through #2 (vt100) (PG)
clear_all_tabs,	tbc	ct	Clear all tab stops (P)
clear_screen,	clear	cl	Clear screen and home cursor (P*)
clr_eol,	el	ee	Clear to end of line (P)
clr_eos,	ed	ed	Clear to end of display (P*)

## TERMINFO(4)

column_address,	hpa	ch	Set cursor column (PG)
command_character,	cmdch	CC	Term. settable cmd char in prototype
cursor_address,	cup	cm	Screen rel. cursor motion row #1 col #2 (PG)
cursor_down,	cu dl	do	Down one line
cursor_home,	home	ho	Home cursor (if no cup)
cursor_invisible,	civis	vi	Make cursor invisible
cursor_left,	cubl	le	Move cursor left one space
cursor_mem_address,	mrcup	CM	Memory relative cursor addressing
cursor_normal,	cnorm	ve	Make cursor appear normal (undo vs/vi)
cursor_right,	cuf l	nd	Non-destructive space (cursor right)
cursor_to_ll,	ll	ll	Last line, first column (if no cup)
cursor_up,	cuul	up	Upline (cursor up)
cursor_visible,	cvvis	vs	Make cursor very visible
delete_character,	dchl	dc	Delete character (P*)
delete_line,	dll	dl	Delete line (P*)
dis_status_line,	dsl	ds	Disable status line
down_half_line,	hd	hd	Half-line down (forward 1/2 linefeed)
enter_alt_charset_mode,	smacs	as	Start alternate character set (P)
enter_blink_mode,	blink	mb	Turn on blinking
enter_bold_mode,	bold	md	Turn on bold (extra bright) mode
enter_ca_mode,	smcup	ti	String to begin programs that use cup
enter_delete_mode,	smdc	dm	Delete mode (enter)
enter_dim_mode,	dim	mh	Turn on half-bright mode
enter_insert_mode,	smir	im	Insert mode (enter);
enter_protected_mode,	prot	mp	Turn on protected mode
enter_reverse_mode,	rev	mr	Turn on reverse video mode
enter_secure_mode,	invis	mk	Turn on blank mode (chars invisible)
enter_standout_mode,	smso	so	Begin stand-out mode
enter_underline_mode,	smul	us	Start underscore mode
erase_chars	ech	ec	Erase #1 characters (PG)
exit_alt_charset_mode,	rmacs	ae	End alternate character set (P)
exit_attribute_mode,	sgr0	me	Turn off all attributes
exit_ca_mode,	rmcup	te	String to end programs that use cup
exit_delete_mode,	rmdc	ed	End delete mode
exit_insert_mode,	rmir	ei	End insert mode
exit_standout_mode,	rmso	se	End stand-out mode
exit_underline_mode,	rmul	ue	End underscore mode
flash_screen,	flash	vb	Visible bell (may not move cursor)
form_feed,	ff	ff	Hardcopy terminal page eject (P*)
from_status_line,	fsl	fs	Return from status line
init_1string,	isl	i1	Terminal initialization string
init_2string,	is2	i2	Terminal initialization string
init_3string,	is3	i3	Terminal initialization string
init_file,	if	if	Name of file containing is
insert_character,	ichl	ic	Insert character (P)
insert_line,	ill	al	Add new blank line (P*)
insert_padding,	ip	ip	Insert pad after character inserted (p*)
key_backspace,	kbs	kb	Sent by backspace key
key_catab,	ktbc	ka	Sent by clear-all-tabs key
key_clear,	kclr	kC	Sent by clear screen or erase key
key_ctab,	kctab	kt	Sent by clear-tab key
key_dc,	kdchl	kD	Sent by delete character key
key_dl,	kdll	kL	Sent by delete line key

key_down,	kcudl	kd	Sent by terminal down arrow key
key_eic,	krmir	kM	Sent by rmir or smir in insert mode
key_col,	kel	kE	Sent by clear-to-end-of-line key
key_eos,	ked	kS	Sent by clear-to-end-of-screen key
key_f0,	kf0	k0	Sent by function key f0
key_f1,	kf1	k1	Sent by function key f1
key_f10,	kf10	ka	Sent by function key f1 0
key_f2,	kf2	k2	Sent by function key f2
key_f3,	kf3	k3	Sent by function key f3
key_f4,	kf4	k4	Sent by function key f4
key_f5,	kf5	k5	Sent by function key f5
key_f6,	kf6	k6	Sent by function key f6
key_f7,	kf7	k7	Sent by function key f7
key_f8,	kf8	k8	Sent by function key f8
key_f9,	kf9	k9	Sent by function key f9
key_home,	khome	kh	Sent by home key
key_ic,	kichl	kI	Sent by ins char/enter ins mode key
key_il,	kill	kA	Sent by insert line
key_left,	kcubl	kl	Sent by terminal left arrow key
key_ll,	kll	kH	Sent by home-down key
key_npage,	knp	kN	Sent by next-page key
key_ppage,	kpp	kP	Sent by previous-page key
key_right,	kcufl	kr	Sent by terminal right arrow key
key_sf,	kind	kF	Sent by scroll-forward/down key
key_sr,	kri	kR	Sent by scroll-backward/up key
key_stab,	khts	kT	Sent by set-tab key
key_up,	kcuul	ku	Sent by terminal up arrow key
keypad_local,	rmkx	ke	Out of "keypad transmit" mode
keypad_xmit,	smkx	ks	Put terminal in "keypad transmit" mode
lab_f0,	lf0	l0	Labels on function key f0 if not f0
lab_f1,	lf1	l1	Labels on function key f1 if not f1
lab_f10,	lf10	la	Labels on function key f10 if not f10
lab_f2,	lf2	l2	Labels on function key f2 if not f2
lab_f3,	lf3	l3	Labels on function key f3 if not f3
lab_f4,	lf4	l4	Labels on function key f4 if not f4
lab_f5,	lf5	l5	Labels on function key f5 if not f5
lab_f6,	lf6	l6	Labels on function key f6 if not f6
lab_f7,	lf7	l7	Labels on function key f7 if not f7
lab_f8,	lf8	l8	Labels on function key f8 if not f8
lab_f9,	lf9	l9	Labels on function key f9 if not f9
meta_on,	smm	mm	Turn on "meta mode" (8th bit)
meta_off,	rmm	mo	Turn off "meta mode"
newline,	nel	nw	New-line (behaves like cr followed by lf)
pad_char,	pad	pc	Pad character (rather than null)
parm_dch,	dch	DC	Delete #1 chars (PG*)
parm_delete_line,	dl	DL	Delete #1 lines (PG*)
parm_down_cursor,	cud	DO	Move cursor down #1 lines (PG*)
parm_ich,	ich	IC	Insert #1 blank chars (PG*)
parm_index,	indn	SF	Scroll forward #1 lines (PG)
parm_insert_line,	il	AL	Add #1 new blank lines (PG*)
parm_left_cursor,	cub	LE	Move cursor left #1 spaces (PG)
parm_right_cursor,	cuf	RI	Move cursor right #1 spaces (PG*)
parm_rindex,	rin	SR	Scroll backward #1 lines (PG)
parm_up_cursor,	cuu	UP	Move cursor up #1 lines (PG*)

## TERMINFO(4)

pkey_key,	pfkey	pk	Prog funct key #1 to type string #2
pkey_local,	pfloc	pl	Prog funct key #1 to execute string #2
pkey_xmit,	pfx	px	Prog funct key #1 to xmit string #2
print_screen,	mc0	ps	Print contents of the screen
prtr_off,	mc4	pf	Turn off the printer
prtr_on,	mc5	po	Turn on the printer
repeat_char,	rep	rp	Repeat char #1 #2 times. (PG*)
reset_lstring,	rs1	r1	Reset terminal completely to sane modes.
reset_2string,	rs2	r2	Reset terminal completely to sane modes.
reset_3string,	rs3	r3	Reset terminal completely to sane modes.
reset_file,	rf	rf	Name of file containing reset string
restore_cursor,	rc	rc	Restore cursor to position of last sc
row_address,	vpa	cv	Vertical position absolute (set row) (PG)
save_cursor,	sc	sc	Save cursor position (P)
scroll_forward,	ind	sf	Scroll text up (P)
scroll_reverse,	ri	sr	Scroll text down (P)
set_attributes,	sgr	sa	Define the video attributes (PG9)
set_tab,	hts	st	Set a tab in all rows, current column
set_window,	wind	wi	Current window is lines #1-#2 cols #3-#4
tab,	ht	ta	Tab to next 8 space hardware tab stop
to_status_line,	tsl	ts	Go to status line, column #1
underline_char,	uc	uc	Underscore one char and move past it
up_half_line,	hu	hu	Half-line up (reverse 1/2 linefeed)
init_prog,	iprog	iP	Path name of program for init
key_a1,	ka1	K1	Upper left of keypad
key_a3,	ka3	K3	Upper right of keypad
key_b2,	kb2	K2	Center of keypad
key_c1,	kcl	K4	Lower left of keypad
key_c3,	kc3	K5	Lower right of keypad
prtr_non,	mc5p	pO	Turn on the printer for #1 bytes

### A Sample Entry

The following entry, which describes the Concept-100, is among the more complex entries in the *terminfo* file as of this writing.

```
concept100 |c100|concept|c104|c100-4p|concept 100,
am, bel=^G, blank=^EH, blink=^EC, clear=^L$<2+>, cnorm=^EW,
cols#80, cr=^M$<9>, cub1=^H, cud1=^J, cuf1=^E=,
cup=^Ea%p1% ' %+%c%p2% ' %+%c,
cuu1=^E; , cvvis=^EW, db, dch1=^E^A$<16+>, dim=^EE, d11=^E^B$<3+>,
ed=^E^C$<16+>, el=^E^U$<16>, eo, flash=^Ek$<20>\EK, ht=^t$<8>,
il1=^E^R$<3+>, in, ind=^J, .ind=^J$<9>, yp=$<16+>,
is2=^EU^Ef^E7^E5^E8^E1^ENH^EK^E200^Eo&^200^Eo^47^E,
kbs=^h, kcub1=^E>, kcucl=^E<, kcufl=^E=, kcuu1=^E; ,
kf1=^E5, kf2=^E6, kf3=^E7, khome=^E7,
lines#24, mir, pb#9600, prot=^EI, rep=^Er%p1%c%p2% ' %+%c$<.2+>,
rev=^ED, rmcup=^Ev $<6>^Ep\r\n, rmir=^E\200, rmkx=^Ex,
rmso=^Ed^Ee, rmul=^Eg, rmul=^Eg, sgr0=^EN\200,
smcup=^EU^Ev 8p^Ep\r, smir=^E^P, smkx=^EX, smso=^EE^ED,
smul=^EG, tabs, ul, vt#8, xenl,
```

Entries may continue onto multiple lines by placing white space at the beginning of each line except the first. Comments may be included on lines beginning with "#". Capabilities in *terminfo* are of three types: Boolean capabilities which indicate that the terminal has some particular feature, numeric capabilities giving the size of the terminal or the size of particular delays, and

string capabilities, which give a sequence which can be used to perform particular terminal operations.

### Types of Capabilities

All capabilities have names. For instance, the fact that the Concept has *automatic margins* (i.e., an automatic return and linefeed when the end of a line is reached) is indicated by the capability **am**. Hence the description of the Concept includes **am**. Numeric capabilities are followed by the character '#' and then the value. Thus **cols**, which indicates the number of columns the terminal has, gives the value '80' for the Concept.

Finally, string valued capabilities, such as **el** (clear to end of line sequence) are given by the two-character code, an '=', and then a string ending at the next following ','. A delay in milliseconds may appear anywhere in such a capability, enclosed in \$<..> brackets, as in **el**=\EK\$<3>, and padding characters are supplied by *tputs* to provide this delay. The delay can be either a number, e.g., '20', or a number followed by an '\*', i.e., '3\*'. A '\*' indicates that the padding required is proportional to the number of lines affected by the operation, and the amount given is the per-affected-unit padding required. (In the case of insert character, the factor is still the number of *lines* affected. This is always one unless the terminal has **xeml** and the software uses it.) When a '\*' is specified, it is sometimes useful to give a delay of the form '3.5' to specify a delay per unit to tenths of milliseconds. (Only one decimal place is allowed.)

A number of escape sequences are provided in the string valued capabilities for easy encoding of characters there. Both **\E** and **\e** map to an ESCAPE character, **\x** maps to a control-x for any appropriate x, and the sequences **\n** **\l** **\r** **\t** **\b** **\f** **\s** give a new-line, linefeed, return, tab, backspace, formfeed, and space. Other escapes include **\^** for '^', **\|** for '|', **\,** for comma, **\:** for ':', and **\0** for null. (**\0** will produce **\200**, which does not terminate a string but behaves as a null character on most terminals.) Finally, characters may be given as three octal digits after a **\**.

Sometimes individual capabilities must be commented out. To do this, put a period before the capability name. For example, see the second **ind** in the example above.

### Preparing Descriptions

We now outline how to prepare descriptions of terminals. The most effective way to prepare a terminal description is by imitating the description of a similar terminal in *terminfo* and to build up a description gradually, using partial descriptions with *vi* to check that they are correct. Be aware that a very unusual terminal may expose deficiencies in the ability of the *terminfo* file to describe it or bugs in *vi*. To easily test a new terminal description you can set the environment variable **TERMINFO** to a path name of a directory containing the compiled description you are working on and programs will look there rather than in *usr/lib/terminfo*. To get the padding for insert line right (if the terminal manufacturer did not document it) a severe test is to edit */etc/passwd* at 9600 baud, delete 16 or so lines from the middle of the screen, then hit the 'u' key several times quickly. If the terminal messes up, more padding is usually needed. A similar test can be used for insert character.

### Basic Capabilities

The number of columns on each line for the terminal is given by the **cols** numeric capability. If the terminal is a CRT, then the number of lines on the screen is given by the **lines** capability. If the terminal wraps around to the beginning of the next line when it reaches the right margin, then it should have the **am** capability. If the terminal can clear its screen, leaving the cursor in the home position, then this is given by the **clear** string capability. If the terminal

overstrikes (rather than clearing a position when a character is struck over) then it should have the `os` capability: If the terminal is a printing terminal, with no soft copy unit, give it both `hc` and `os`. (`os` applies to storage scope terminals, such as TEKTRONIX 4010 series, as well as hard copy and APL terminals.) If there is a code to move the cursor to the left edge of the current row, give this as `cr`. (Normally this will be carriage return, control M.) If there is a code to produce an audible signal (bell, beep, etc) give this as `bel`.

If there is a code to move the cursor one position to the left (such as backspace) that capability should be given as `cub1`. Similarly, codes to move to the right, up, and down should be given as `cuf1`, `cuu1`, and `cod1`. These local cursor motions should not alter the text they pass over, for example, you would not normally use `'cuf1='` because the space would erase the character moved over.

A very important point here is that the local cursor motions encoded in *terminfo* are undefined at the left and top edges of a CRT terminal. Programs should never attempt to backspace around the left edge, unless `bw` is given, and never attempt to go up locally off the top. In order to scroll text up, a program will go to the bottom left corner of the screen and send the `ind` (index) string.

To scroll text down, a program goes to the top left corner of the screen and sends the `ri` (reverse index) string. The strings `ind` and `ri` are undefined when not on their respective corners of the screen.

Parameterized versions of the scrolling sequences are `indn` and `rin` which have the same semantics as `ind` and `ri` except that they take one parameter, and scroll that many lines. They are also undefined except at the appropriate edge of the screen.

The `am` capability tells whether the cursor sticks at the right edge of the screen when text is output, but this does not necessarily apply to a `cuf1` from the last column. The only local motion which is defined from the left edge is if `bw` is given, then a `cub1` from the left edge will move to the right edge of the previous row. If `bw` is not given, the effect is undefined. This is useful for drawing a box around the edge of the screen, for example. If the terminal has switch selectable automatic margins, the *terminfo* file usually assumes that this is on; i.e., `am`. If the terminal has a command which moves to the first column of the next line, that command can be given as `nel` (new-line). It does not matter if the command clears the remainder of the current line, so if the terminal has no `cr` and if it may still be possible to craft a working `nel` out of one or both of them.

These capabilities suffice to describe hardcopy and glass-tty terminals. Thus the model 33 TELETYPE is described as

```
33|tty33|tty|model 33 TELETYPE,
bel=^G, cols#72, cr=^M, cud1=^J, hc, ind=^J, os,
```

while the Lear Siegler ADM-3 is described as

```
adm3|3|lsi adm3,
am, bel=^G, clear=^Z, cols#80, cr=^M, cub1=^H, cud1=^J,
ind=^J, lines#24,
```

### Parameterized Strings

Cursor addressing and other strings requiring parameters in the terminal are described by a parameterized string capability, with *printf*(3S) like escapes `%x` in it. For example, to address the cursor, the `cup` capability is given, using two parameters: the row and column to address to. (Rows and columns are numbered from zero and refer to the physical screen visible to the user, not to any unseen memory.) If the terminal has memory relative cursor addressing, that can be indicated by `mrcup`.

The parameter mechanism uses a stack and special % codes to manipulate it. Typically a sequence will push one of the parameters onto the stack and then print it in some format. Often more complex operations are necessary.

The % encodings have the following meanings:

%%	outputs '%'
%d	print pop() as in printf
%2d	print pop() like %2d
%3d	print pop() like %3d
%02d	
%03d	as in printf
%c	print pop() gives %c
%s	print pop() gives %s
%p[1-9]	push ith parm
%P[a-z]	set variable [a-z] to pop()
%g[a-z]	get variable [a-z] and push it
%'c'	char constant c
%(nn)	integer constant nn
%+ %- %* %/ %m	arithmetic (%m is mod): push(pop() op pop())
%& %  %^	bit operations: push(pop() op pop())
%= %> %<	logical operations: push(pop() op pop())
%! %'	unary operations push(op pop())
%i	add 1 to first two parms (for ANSI terminals)
%? expr %t thenpart %e elsepart %;	if-then-else, %e elsepart is optional.
	else-if's are possible ala Algol 68:
%? c <sub>1</sub> %t b <sub>1</sub> %e c <sub>2</sub> %t b <sub>2</sub> %e c <sub>3</sub> %t b <sub>3</sub> %e c <sub>4</sub> %t b <sub>4</sub> %e %;	c <sub>i</sub> are conditions, b <sub>i</sub> are bodies.

Binary operations are in postfix form with the operands in the usual order. That is, to get x-5 one would use "%gx%(5)%-".

Consider the HP2645, which, to get to row 3 and column 12, needs to be sent `\E&a12c03Y` padded for 6 milliseconds. Note that the order of the rows and columns is inverted here, and that the row and column are printed as two digits. Thus its `cup` capability is `cup=6\E&%p2%2dc%p1%2dY`.

The Microterm ACT-IV needs the current row and column sent preceded by a `^T`, with the row and column simply encoded in binary, `cup=~T%p1%c%p2%c`. Terminals which use %c need to be able to backspace the cursor (`cuB1`), and to move the cursor up one line on the screen (`cuu1`). This is necessary because it is not always safe to transmit `\n ^D` and `\r`, as the system may change or discard them. (The library routines dealing with terminfo set tty modes so that tabs are never expanded, so `\t` is safe to send. This turns out to be essential for the Ann Arbor 4080.)

A final example is the LSI ADM-3a, which uses row and column offset by a blank character, thus `cup=\E=%p1%' '%+%c%p2%' '%+%c`. After sending `\E=`, this pushes the first parameter, pushes the ASCII value for a space (32), adds them (pushing the sum on the stack in place of the two previous values) and outputs that value as a character. Then the same is done for the second parameter. More complex arithmetic is possible using the stack.

If the terminal has row or column absolute cursor addressing, these can be given as single parameter capabilities `hpa` (horizontal position absolute) and `vpa` (vertical position absolute). Sometimes these are shorter than the more general

## TERMINFO(4)

two parameter sequence (as with the hp2645) and can be used in preference to `cup`. If there are parameterized local motions (e.g., move  $n$  spaces to the right) these can be given as `cud`, `cub`, `cuf`, and `cuu` with a single parameter indicating how many spaces to move. These are primarily useful if the terminal does not have `cup`, such as the TEKTRONIX 4025.

### Cursor Motions

If the terminal has a fast way to home the cursor (to very upper left corner of screen) then this can be given as `home`; similarly a fast way of getting to the lower left-hand corner can be given as `ll`; this may involve going up with `cuul` from the home position, but a program should never do this itself (unless `ll` does) because it can make no assumption about the effect of moving up from the home position. Note that the home position is the same as addressing to (0,0): to the top left corner of the screen, not of memory. (Thus, the `\EH` sequence on HP terminals cannot be used for `home`.)

### Area Clears

If the terminal can clear from the current position to the end of the line, leaving the cursor where it is, this should be given as `el`. If the terminal can clear from the current position to the end of the display, then this should be given as `ed`. `Ed` is only defined from the first column of a line. (Thus, it can be simulated by a request to delete a large number of lines, if a true `ed` is not available.)

### Insert/Delete Line

If the terminal can open a new blank line before the line where the cursor is, this should be given as `ill`; this is done only from the first position of a line. The cursor must then appear on the newly blank line. If the terminal can delete the line which the cursor is on, then this should be given as `dll`; this is done only from the first position on the line to be deleted. Versions of `ill` and `dll` which take a single parameter and insert or delete that many lines can be given as `il` and `dl`. If the terminal has a settable scrolling region (like the vt100) the command to set this can be described with the `csr` capability, which takes two parameters: the top and bottom lines of the scrolling region. The cursor position is, alas, undefined after using this command. It is possible to get the effect of insert or delete line using this `-` command; the `sc` and `rc` (save and restore cursor) commands are also useful. Inserting lines at the top or bottom of the screen can also be done using `ri` or `ind` on many terminals without a true insert/delete line, and is often faster even on terminals with those features.

If the terminal has the ability to define a window as part of memory, which all commands affect, it should be given as the parameterized string `wind`. The four parameters are the starting and ending lines in memory and the starting and ending columns in memory, in that order.

If the terminal can retain display memory above, then the `da` capability should be given; if display memory can be retained below, then `db` should be given. These indicate that deleting a line or scrolling may bring non-blank lines up from below or that scrolling back with `ri` may bring down non-blank lines.

### Insert/Delete Character

There are two basic kinds of intelligent terminals with respect to insert/delete characters which can be described using `terminfo`. The most common insert/delete character operations affect only the characters on the current line and shift characters off the end of the line rigidly. Other terminals, such as the Concept 100 and the Perkin Elmer Owl, make a distinction between typed and untyped blanks on the screen, shifting upon an insert or delete only to an untyped blank on the screen which is either eliminated, or expanded to two untyped blanks. You can determine the kind of terminal you have by clearing



the screen and then typing text separated by cursor motions. Type `abc def` using local cursor motions (not spaces) between the `abc` and the `def`. Then position the cursor before the `abc` and put the terminal in insert mode. If typing characters causes the rest of the line to shift rigidly and characters to fall off the end, then your terminal does not distinguish between blanks and untyped positions. If the `abc` shifts over to the `def` which then move together around the end of the current line and onto the next as you insert, you have the second type of terminal, and should give the capability `in`, which stands for insert null. While these are two logically separate attributes (one line vs. multiline insert mode, and special treatment of untyped spaces) we have seen no terminals whose insert mode cannot be described with the single attribute.

Terminfo can describe both terminals which have an insert mode, and terminals which send a simple sequence to open a blank position on the current line. Give as `smir` the sequence to get into insert mode. Give as `rmir` the sequence to leave insert mode. Now give as `ichl` any sequence needed to be sent just before sending the character to be inserted. Most terminals with a true insert mode will not give `ichl`; terminals which send a sequence to open a screen position should give it here. (If your terminal has both, insert mode is usually preferable to `ichl`. Do not give both unless the terminal actually requires both to be used in combination.) If post insert padding is needed, give this as a number of milliseconds in `ip` (a string option). Any other sequence which may need to be sent after an insert of a single character may also be given in `ip`. If your terminal needs both to be placed into an 'insert mode' and a special code to precede each inserted character, then both `smir/rmir` and `ichl` can be given, and both will be used. The `ich` capability, with one parameter, `n`, will repeat the effects of `ichl n` times.

It is occasionally necessary to move around while in insert mode to delete characters on the same line (e.g., if there is a tab after the insertion position). If your terminal allows motion while in insert mode you can give the capability `mir` to speed up inserting in this case. Omitting `mir` will affect only speed. Some terminals (notably Datamedia's) must not have `mir` because of the way their insert mode works.

Finally, you can specify `dchl` to delete a single character, `dch` with one parameter, `n`, to delete `n` characters, and delete mode by giving `smdc` and `rmdc` to enter and exit delete mode (any mode the terminal needs to be placed in for `dchl` to work).

A command to erase `n` characters (equivalent to outputting `n` blanks without moving the cursor) can be given as `ech` with one parameter.

### Highlighting, Underlining, and Visible Bells

If your terminal has one or more kinds of display attributes, these can be represented in a number of different ways. You should choose one display form as *standout mode*, representing a good, high contrast, easy-on-the-eyes, format for highlighting error messages and other attention getters. (If you have a choice, reverse video plus half-bright is good, or reverse video alone.) The sequences to enter and exit standout mode are given as `smso` and `rmso`, respectively. If the code to change into or out of standout mode leaves one or even two blank spaces on the screen, as the TVI 912 and Teleray 1061 do, then `xmc` should be given to tell how many spaces are left.

Codes to begin underlining and end underlining can be given as `smul` and `rmul` respectively. If the terminal has a code to underline the current character and move the cursor one space to the right, such as the Microterm Mime, this can be given as `uc`.

## TERMINFO(4)

Other capabilities to enter various highlighting modes include **blink** (blinking) **bold** (bold or extra bright) **dim** (dim or half-bright) **invis** (blinking or invisible text) **prot** (protected) **rev** (reverse video) **sgr0** (turn off *all* attribute modes) **smacs** (enter alternate character set mode) and **rmacs** (exit alternate character set mode). Turning on any of these modes singly may or may not turn off other modes.

If there is a sequence to set arbitrary combinations of modes, this should be given as **sgr** (set attributes), taking 9 parameters. Each parameter is either 0 or 1, as the corresponding attribute is on or off. The 9 parameters are, in order: standout, underline, reverse, blink, dim, bold, blank, protect, alternate character set. Not all modes need be supported by **sgr**, only those for which corresponding separate attribute commands exist.

Terminals with the "magic cookie" glitch (**xmc**) deposit special "cookies" when they receive mode-setting sequences, which affect the display algorithm rather than having extra bits for each character. Some terminals, such as the HP 2621, automatically leave standout mode when they move to a new line or the cursor is addressed. Programs using standout mode should exit standout mode before moving the cursor or sending a new-line, unless the **msgr** capability, asserting that it is safe to move in standout mode, is present.

If the terminal has a way of flashing the screen to indicate an error quietly (a bell replacement) then this can be given as **flash**; it must not move the cursor.

If the cursor needs to be made more visible than normal when it is not on the bottom line (to make, for example, a non-blinking underline into an easier to find block or blinking underline) give this sequence as **cvvis**. If there is a way to make the cursor completely invisible, give that as **civis**. The capability **cnorm** should be given which undoes the effects of both of these modes.

If the terminal needs to be in a special mode when running a program that uses these capabilities, the codes to enter and exit this mode can be given as **smcup** and **rmcup**. This arises, for example, from terminals like the Concept with more than one page of memory. If the terminal has only memory relative cursor addressing and not screen relative cursor addressing, a one screen-sized window must be fixed into the terminal for cursor addressing to work properly. This is also used for the TEKTRONIX 4025, where **smcup** sets the command character to be the one used by **terminfo**.

If your terminal correctly generates underlined characters (with no special codes needed) even though it does not overstrike, then you should give the capability **ul**. If overstrikes are erasable with a blank, then this should be indicated by giving **eo**.

### Keypad

If the terminal has a keypad that transmits codes when the keys are pressed, this information can be given. Note that it is not possible to handle terminals where the keypad only works in local (this applies, for example, to the unshifted HP 2621 keys). If the keypad can be set to transmit or not transmit, give these codes as **smkx** and **rmkx**. Otherwise the keypad is assumed to always transmit. The codes sent by the left arrow, right arrow, up arrow, down arrow, and home keys can be given as **kcub1**, **kcuf1**, **kcuu1**, **kcud1**, and **khome** respectively. If there are function keys such as **f0**, **f1**, ..., **f10**, the codes they send can be given as **kf0**, **kf1**, ..., **kf10**. If these keys have labels other than the default **f0** through **f10**, the labels can be given as **lf0**, **lf1**, ..., **lf10**. The codes transmitted by certain other special keys can be given: **kll** (home down), **kbs** (backspace), **ktbc** (clear all tabs), **kctab** (clear the tab stop in this column), **kclr** (clear screen or erase key), **kdcbl** (delete character), **kdl1** (delete line), **krmir** (exit insert mode), **kel** (clear to end of line), **ked** (clear to end of screen),

**ki**cb1 (insert character or enter insert mode), **kill** (insert line), **knp** (next page), **kpp** (previous page), **kind** (scroll forward/down), **kri** (scroll backward/up), **kh**ts (set a tab stop in this column). In addition, if the keypad has a 3-by-3 array of keys including the four arrow keys, the other five keys can be given as **ka**1, **ka**3, **kb**2, **kcl**, and **kc**3. These keys are useful when the effects of a 3-by-3 directional pad are needed.

### Tabs and Initialization

If the terminal has hardware tabs, the command to advance to the next tab stop can be given as **ht** (usually control I). A "backtab" command which moves leftward to the next tab stop can be given as **cbt**. By convention, if the TELETYPE modes indicate that tabs are being expanded by the computer rather than being sent to the terminal, programs should not use **ht** or **cbt** even if they are present, since the user may not have the tab stops properly set. If the terminal has hardware tabs which are initially set every *n* spaces when the terminal is powered up, the numeric parameter **it** is given, showing the number of spaces the tabs are set to. This is normally used by the *tset* command to determine whether to set the mode for hardware tab expansion, and whether to set the tab stops. If the terminal has tab stops that can be saved in nonvolatile memory, the terminfo description can assume that they are properly set.

Other capabilities include **is**1, **is**2, and **is**3, initialization strings for the terminal, **ipro**g, the path name of a program to be run to initialize the terminal, and **if**, the name of a file containing long initialization strings. These strings are expected to set the terminal into modes consistent with the rest of the terminfo description. They are normally sent to the terminal, by the *tset* program, each time the user logs in. They will be printed in the following order: **is**1; **is**2; setting tabs using **tbc** and **ht**s; **if**; running the program **ipro**g; and finally **is**3. Most initialization is done with **is**2. Special terminal modes can be set up without duplicating strings by putting the common sequences in **is**2 and special cases in **is**1 and **is**3. A pair of sequences that does a harder reset from a totally unknown state can be analogously given as **rs**1, **rs**2, **rf**, and **rs**3, analogous to **is**2 and **if**. These strings are output by the *reset* program, which is used when the terminal gets into a wedged state. Commands are normally placed in **rs**2 and **rf** only if they produce annoying effects on the screen and are not necessary when logging in. For example, the command to set the vt100 into 80-column mode would normally be part of **is**2, but it causes an annoying glitch of the screen and is not normally needed since the terminal is usually already in 80-column mode.

If there are commands to set and clear tab stops, they can be given as **tbc** (clear all tab stops) and **ht**s (set a tab stop in the current column of every row). If a more complex sequence is needed to set the tabs than can be described by this, the sequence can be placed in **is**2 or **if**.

### Delays

Certain capabilities control padding in the TELETYPE driver. These are primarily needed by hard copy terminals, and are used by the *tset* program to set TELETYPE modes appropriately. Delays embedded in the capabilities **cr**, **ind**, **cub**1, **ff**, and **tab** will cause the appropriate delay bits to be set in the TELETYPE driver. If **pb** (padding baud rate) is given, these values can be ignored at baud rates below the value of **ph**.

### Miscellaneous

If the terminal requires other than a null (zero) character as a pad, then this can be given as **pad**. Only the first character of the **pad** string is used.

If the terminal has an extra "status line" that is not normally used by software, this fact can be indicated. If the status line is viewed as an extra line below

## TERMINFO(4)

the bottom line, into which one can cursor address normally (such as the Heathkit h19's 25th line, or the 24th line of a vt100 which is set to a 23-line scrolling region), the capability **hs** should be given. Special strings to go to the beginning of the status line and to return from the status line can be given as **tsl** and **fsl**. (**fsl** must leave the cursor position in the same place it was before **tsl**. If necessary, the **sc** and **rc** strings can be included in **tsl** and **fsl** to get this effect.) The parameter **tsl** takes one parameter, which is the column number of the status line the cursor is to be moved to. If escape sequences and other special commands, such as **tab**, work while in the status line, the flag **eslok** can be given. A string which turns off the status line (or otherwise erases its contents) should be given as **dsl**. If the terminal has commands to save and restore the position of the cursor, give them as **sc** and **rc**. The status line is normally assumed to be the same width as the rest of the screen, e.g., **cols**. If the status line is a different width (possibly because the terminal does not allow an entire line to be loaded) the width, in columns, can be indicated with the numeric parameter **wsl**.

If the terminal can move up or down half a line, this can be indicated with **hu** (half-line up) and **hd** (half-line down). This is primarily useful for superscripts and subscripts on hardcopy terminals. If a hardcopy terminal can eject to the next page (form feed), give this as **ff** (usually control L).

If there is a command to repeat a given character a given number of times (to save time transmitting a large number of identical characters) this can be indicated with the parameterized string **rep**. The first parameter is the character to be repeated and the second is the number of times to repeat it. Thus, **tparam(repeat\_char, 'x', 10)** is the same as 'xxxxxxxx'.

If the terminal has a settable command character, such as the TEKTRONIX 4025, this can be indicated with **cmdch**. A prototype command character is chosen which is used in all capabilities. This character is given in the **cmdch** capability to identify it. The following convention is supported on some UNIX systems: The environment is to be searched for a **CC** variable, and if found, all occurrences of the prototype character are replaced with the character in the environment variable.

Terminal descriptions that do not represent a specific kind of known terminal, such as *switch*, *dialup*, *patch*, and *network*, should include the **gn** (generic) capability so that programs can complain that they do not know how to talk to the terminal. (This capability does not apply to *virtual* terminal descriptions for which the escape sequences are known.)

If the terminal uses **xon/xoff** handshaking for flow control, give **xon**. Padding information should still be included so that routines can make better decisions about costs, but actual pad characters will not be transmitted.

If the terminal has a "meta key" which acts as a shift key, setting the 8th bit of any character transmitted, this fact can be indicated with **km**. Otherwise, software will assume that the 8th bit is parity and it will usually be cleared. If strings exist to turn this "meta mode" on and off, they can be given as **smm** and **rmm**.

If the terminal has more lines of memory than will fit on the screen at once, the number of lines of memory can be indicated with **lm**. A value of **lm#0** indicates that the number of lines is not fixed, but that there is still more memory than fits on the screen.

If the terminal is one of those supported by the UNIX system virtual terminal protocol, the terminal number can be given as **vt**.

Media copy strings which control an auxiliary printer connected to the terminal can be given as **mc0**: print the contents of the screen, **mc4**: turn off the

printer, and **mc5**: turn on the printer. When the printer is on, all text sent to the terminal will be sent to the printer. It is undefined whether the text is also displayed on the terminal screen when the printer is on. A variation **mc5p** takes one parameter, and leaves the printer on for as many characters as the value of the parameter, then turns the printer off. The parameter should not exceed 255. All text, including **mc4**, is transparently passed to the printer while an **mc5p** is in effect.

Strings to program function keys can be given as **pfkey**, **pfloc**, and **pfx**. Each of these strings takes two parameters: the function key number to program (from 0 to 10) and the string to program it with. Function key numbers out of this range may program undefined keys in a terminal-dependent manner. The difference between the capabilities is that **pfkey** causes pressing the given key to be the same as the user typing the given string; **pfloc** causes the string to be executed by the terminal in local; and **pfx** causes the string to be transmitted to the computer.

### Glitches and Braindamage

Hazeltine terminals, which do not allow “” characters to be displayed should indicate **bz**.

Terminals which ignore a linefeed immediately after an **am** wrap, such as the Concept and vt100, should indicate **xenl**.

If **el** is required to get rid of standout (instead of merely writing normal text on top of it), **xhp** should be given.

Teleray terminals, where tabs turn all characters moved over to blanks, should indicate **xt** (destructive tabs). This glitch is also taken to mean that it is not possible to position the cursor on top of a “magic cookie”, that to erase standout mode it is instead necessary to use delete and insert line.

The Beehive Superbee, which is unable to correctly transmit the escape or control C characters, has **xsbs**, indicating that the **f1** key is used for escape and **f2** for control C. (Only certain Superbees have this problem, depending on the ROM.)

Other specific terminal problems may be corrected by adding more capabilities of the form **xx**.

### Similar Terminals

If there are two very similar terminals, one can be defined as being just like the other with certain exceptions. The string capability **use** can be given with the name of the similar terminal. The capabilities given before **use** override those in the terminal type invoked by **use**. A capability can be cancelled by placing **xx@** to the left of the capability definition, where **xx** is the capability. For example, the entry

```
2621-nl, smkx@, rmkx@, use=2621,
```

defines a 2621-nl that does not have the **smkx** or **rmkx** capabilities, and hence does not turn on the function key labels when in visual mode. This is useful for different modes for a terminal, or for different user preferences.

### FILES

`/usr/lib/terminfo/?/*` files containing terminal descriptions

### SEE ALSO

`curses(3X)`, `printf(3S)`, `term(5)`,  
`tic(1M)` in the Runtime System manual.

## UTMP(4)

### NAME

utmp, wtmp - utmp and wtmp entry formats

### SYNOPSIS

```
#include <sys/types.h>
#include <utmp.h>
```

### DESCRIPTION

These files, which hold user and accounting information for such commands as *who(1)*, *write(1)*, and *login(1)*, have the following structure as defined by *<utmp.h>*:

```
#define  UTMP_FILE  "/etc/utmp"
#define  WTMP_FILE  "/etc/wtmp"
#define  ut_name    ut_user

struct utmp {
    char    ut_user[8];           /* User login name */
    char    ut_id[4];           /* /etc/inittab id (usually line #) */
    char    ut_line[12];        /* device name (console, lxxx) */
    short   ut_pid;             /* process id */
    short   ut_type;            /* type of entry */
    struct  exit_status {
        short  e_termination; /* Process termination status */
        short  e_exit;         /* Process exit status */
    } ut_exit;                  /* The exit status of a process
    * marked as DEAD_PROCESS. */
    time_t   ut_time;           /* time entry was made */
};

/* Definitions for ut_type */
#define  EMPTY      0
#define  RUN_LVL    1
#define  BOOT_TIME  2
#define  OLD_TIME   3
#define  NEW_TIME   4
#define  INIT_PROCESS 5           /* Process spawned by "init" */
#define  LOGIN_PROCESS 6         /* A "getty" process waiting for login */
#define  USER_PROCESS 7          /* A user process */
#define  DEAD_PROCESS 8
#define  ACCOUNTING  9
#define  UTMAXTYPE   ACCOUNTING /* Largest legal value of ut_type */

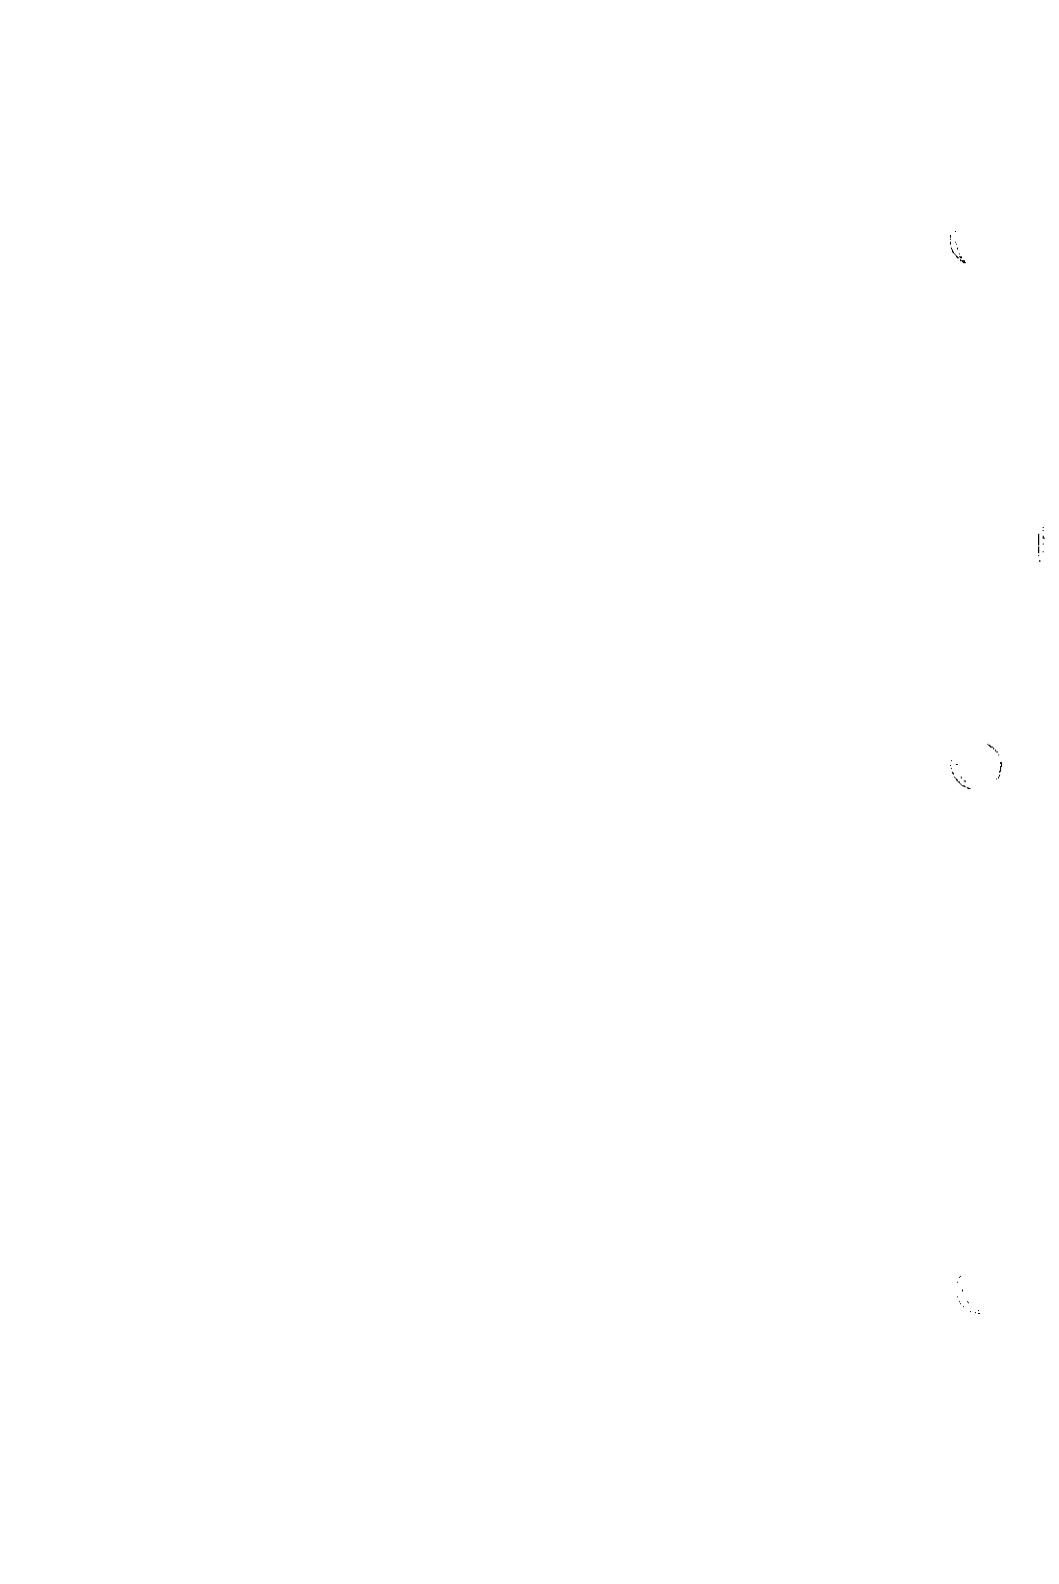
/* Special strings or formats used in the "ut_line" field when */
/* accounting for something other than a process */
/* No string for the ut_line field can be more than 11 chars + */
/* a NULL in length */
#define  RUNLVL_MSG  "run-level %c"
#define  BOOT_MSG   "system boot"
#define  OTIME_MSG  "old time"
#define  NTIME_MSG  "new time"
```

**FILES**

/usr/include/utmp.h  
/etc/utmp  
/etc/wtmp

**SEE ALSO**

getut(3C),  
login(1), who(1), write(1) in the Runtime System manual.





## TABLE OF CONTENTS OF MISCELLANY

### 5. Miscellaneous

intro.....	introduction to miscellany
ascii.....	map of ASCII character set
environ.....	user environment
fcntl.....	file control options
math.....	math functions and constants
prof.....	profile within a function
regexp.....	regular expression cope and match routines
stat.....	data returned by stat system call
term.....	conventional names for terminals
types.....	primitive system data types
values.....	machine-dependent values
varargs.....	handle variable argument list

1

2

3

**NAME**

intro - introduction to miscellany

**DESCRIPTION**

This section describes miscellaneous facilities such as macro packages, character set tables, etc.

S-5

# ASCII(5)

## NAME

ascii - map of ASCII character set

## SYNOPSIS

cat /usr/pub/ascii

## DESCRIPTION

*Ascii* is a map of the ASCII character set, giving both octal and hexadecimal equivalents of each character, to be printed as needed. It contains:

000 nul	001 soh	002 stx	003 etx	004 eot	005 enq	006 ack	007 bel
010 bs	011 ht	012 nl	013 vt	014 np	015 cr	016 so	017 si
020 dle	021 dcl	022 dc2	023 dc3	024 dc4	025 nak	026 syn	027 etb
030 can	031 em	032 sub	033 esc	034 fs	035 gs	036 rs	037 us
040 sp	041 !	042 "	043 #	044 \$	045 %	046 &	047
050 (	051 )	052 *	053 +	054 ,	055 -	056 .	057 /
060 0	061 1	062 2	063 3	064 4	065 5	066 6	067 7
070 8	071 9	072 :	073 ;	074 <	075 =	076 >	077 ?
100 @	101 A	102 B	103 C	104 D	105 E	106 F	107 G
110 H	111 I	112 J	113 K	114 L	115 M	116 N	117 O
120 P	121 Q	122 R	123 S	124 T	125 U	126 V	127 w
130 X	131 Y	132 Z	133 [	134 \	135 ]	136 ^	137 _
140	141 a	142 b	143 c	144 d	145 e	146 f	147 g
150 h	151 i	152 j	153 k	154 l	155 m	156 n	157 o
160 p	161 q	162 r	163 s	164 t	165 u	166 v	167 w
170 x	171 y	172 z	173 {	174	175 }	176 ~	177 del

00 nul	01 soh	02 stx	03 etx	04 eot	05 enq	06 ack	07 bel
08 bs	09 ht	0a nl	0b vt	0c np	0d cr	0e so	0f si
10 dle	11 dcl	12 dc2	13 dc3	14 dc4	15 nak	16 syn	17 etb
18 can	19 em	1a sub	1b esc	1c fs	1d gs	1e rs	1f us
20 sp	21 !	22 "	23 #	24 \$	25 %	26 &	27
28 (	29 )	2a *	2b +	2c ,	2d -	2e .	2f /
30 0	31 1	32 2	33 3	34 4	35 5	36 6	37 7
38 8	39 9	3a :	3b ;	3c <	3d =	3e >	3f ?
40 @	41 A	42 B	43 C	44 D	45 E	46 F	47 G
48 H	49 I	4a J	4b K	4c L	4d M	4e N	4f O
50 P	51 Q	52 R	53 S	54 T	55 U	56 V	57 w
58 X	59 Y	5a Z	5b [	5c \	5d ]	5e ^	5f _
60	61 a	62 b	63 c	64 d	65 e	66 f	67 g
68 h	69 i	6a j	6b k	6c l	6d m	6e n	6f o
70 p	71 q	72 r	73 s	74 t	75 u	76 v	77 w
78 x	79 y	7a z	7b {	7c	7d }	7e ~	7f del

## FILES

/usr/pub/ascii

## NAME

environ - user environment

## DESCRIPTION

An array of strings called the "environment" is made available by *exec(2)* when a process begins. By convention, these strings have the form "name=value". The following names are used by various commands:

**PATH** The sequence of directory prefixes that *sh(1)*, *time(1)*, *nice(1)*, *nohup(1)*, etc., apply in searching for a file known by an incomplete path name. The prefixes are separated by colons (:). *Login(1)* sets **PATH=:/bin:/usr/bin**.

**HOME** Name of the user's login directory, set by *login(1)* from the password file *passwd(4)*.

**TERM** The kind of terminal for which output is to be prepared. This information is used by commands, such as *mm(1)* or *tplot(1G)*, which may exploit special capabilities of that terminal.

**TZ** Time zone information. The format is **xxxnzzz** where **xxx** is standard local time zone abbreviation, **n** is the difference in hours from GMT, and **zzz** is the abbreviation for the daylight-saving local time zone, if any; for example, **EST5EDT**.

Further names may be placed in the environment by the *export* command and "name=value" arguments in *sh(1)*, or by *exec(2)*. It is unwise to conflict with certain shell variables that are frequently exported by **.profile** files: **MAIL**, **PS1**, **PS2**, **IFS**.

## SEE ALSO

*exec(2)*.

*env(1)*, *login(1)*, *sh(1)*, *nice(1)*, *nohup(1)*, *time(1)*, *tplot(1G)* in the Runtime System manual.

*mm(1)* in the "User Reference Manual" chapter of the Text Preparation System manual.

## FCNTL(5)

### NAME

fcntl - file control options

### SYNOPSIS

```
#include <fcntl.h>
```

### DESCRIPTION

The *fcntl(2)* function provides for control over open files. This include file describes *requests* and *arguments* to *fcntl* and *open(2)*.

```
/* Flag values accessible to open(2) and fcntl(2) */
/* (The first three can only be set by open) */
#define O_RDONLY 0
#define O_WRONLY 1
#define O_RDWR 2
#define O_NDELAY 04 /* Nonblocking I/O */
#define O_APPEND 010 /* append (writes guaranteed at the end) */
#define O_SYNC 020 /* synchronous write options */

/* Flag values accessible only to open(2) */
#define O_CREAT 00400 /* open with file create (uses third open arg)*/
#define O_TRUNC 01000 /* open with truncation */
#define O_EXCL 02000 /* exclusive open */

/* fcntl(2) requests */
#define F_DUPFD 0 /* Duplicate files */
#define F_GETFD 1 /* Get files flags */
#define F_SETFD 2 /* Set files flags */
#define F_GETFL 3 /* Get file flags */
#define F_SETFL 4 /* Set file flags */
#define F_GETLK 5 /* Get blocking file locks */
#define F_SETLK 6 /* Set or clear file locks and fail on busy */
#define F_SETLKW 7 /* Set or clear file locks and wait on busy */

/* file segment locking control structure */
struct flock {
    short l_type;
    short l_whence;
    long l_start;
    long l_len; /* if 0 then until EOF */
    int l_pid; /* returned with F_GETLK */
};

/* file segment locking types */
#define F_RDLCK 01 /* Read lock */
#define F_WRLCK 02 /* Write lock */
#define F_UNLCK 03 /* Remove locks */
```

### SEE ALSO

fcntl(2), open(2).

**NAME**  
math — math functions and constants

**SYNOPSIS**  
`#include <math.h>`

**DESCRIPTION**  
This file contains declarations of all the functions in the Math Library (described in Section 3M), as well as various functions in the C Library (Section 3C) that return floating-point values.

It defines the structure and constants used by the *matherr*(3M) error-handling mechanisms, including the following constant used as an error-return value:

**HUGE**           The maximum value of a single-precision floating-point number.

The following mathematical constants are defined for user convenience:

**M\_E**             The base of natural logarithms ( $e$ ).

**M\_LOG2E**       The base-2 logarithm of  $e$ .

**M\_LOG10E**      The base-10 logarithm of  $e$ .

**M\_LN2**          The natural logarithm of 2.

**M\_LN10**         The natural logarithm of 10.

**M\_PI**            The ratio of the circumference of a circle to its diameter. (There are also several fractions of its reciprocal and its square root.)

**M\_SQRT2**        The positive square root of 2.

**M\_SQRT1\_2**      The positive square root of  $1/2$ .

For the definitions of various machine-dependent "constants," see the description of the *<values.h>* header file.

**FILES**  
`/usr/include/math.h`

**SEE ALSO**  
`intro(3)`, `matherr(3M)`, `values(5)`.

## PROF(5)

### NAME

prof — profile within a function

### SYNOPSIS

```
#define MARK
#include <prof.h>
void MARK (name)
```

### DESCRIPTION

*MARK* will introduce a mark called *name* that will be treated the same as a function entry point. Execution of the mark will add to a counter for that mark, and program-counter time spent will be accounted to the immediately preceding mark or to the function if there are no preceding marks within the active function.

*Name* may be any combination of up to six letters, numbers or underscores. Each *name* in a single compilation must be unique, but may be the same as any ordinary program symbol.

For marks to be effective, the symbol *MARK* must be defined before the header file *<prof.h>* is included. This may be defined by a preprocessor directive as in the synopsis, or by a command line argument, i.e:

```
cc -p -DMARK foo.c
```

If *MARK* is not defined, the *MARK(name)* statements may be left in the source files containing them and will be ignored.

### EXAMPLE

In this example, marks can be used to determine how much time is spent in each loop. Unless this example is compiled with *MARK* defined on the command line, the marks are ignored.

```
#include <prof.h>

foo ( )
{
    int i, j;
    .
    .
    .
    MARK(loop1);
    for (i = 0; i < 2000; i++) {
        . . .
    }
    MARK(loop2);
    for (j = 0; j < 2000; j++) {
        . . .
    }
}
```

### SEE ALSO

profil(2), monitor(3C).  
prof(1) in the Runtime System manual.



## NAME

regexp — regular expression compile and match routines

## SYNOPSIS

```
#define INIT <declarations>
#define GETC() <getc code>
#define PEEKC() <peekc code>
#define UNGETC(c) <ungetc code>
#define RETURN(pointer) <return code>
#define ERROR(val) <error code>
#include <regexp.h>

char *compile (instring, expbuf, endbuf, eof)
char *instring, *expbuf, *endbuf;
int eof;

int step (string, expbuf)
char *string, *expbuf;

extern char *loc1, *loc2, *locs;
extern int circf, sed, nbra;
```

## DESCRIPTION

This page describes general-purpose regular expression matching routines in the form of *ed*(1), defined in */usr/include/regexp.h*. Programs such as *ed*(1), *sed*(1), *grep*(1), *bs*(1), *expr*(1), etc., which perform regular expression matching use this source file. In this way, only this file need be changed to maintain regular expression compatibility.

The interface to this file is unpleasantly complex. Programs that include this file must have the following five macros declared before the "#include" <regexp.h>" statement. These macros are used by the *compile* routine.

GETC()	Return the value of the next character in the regular expression pattern. Successive calls to GETC() should return successive characters of the regular expression.
PEEKC()	Return the next character in the regular expression. Successive calls to PEEKC() should return the same character [which should also be the next character returned by GETC()].
UNGETC(c)	Cause the argument <i>c</i> to be returned by the next call to GETC() [and PEEKC()]. No more than one character of pushback is ever needed and this character is guaranteed to be the last character read by GETC(). The value of the macro UNGETC( <i>c</i> ) is always ignored.
RETURN(pointer)	This macro is used on normal exit of the <i>compile</i> routine. The value of the argument <i>pointer</i> is a pointer to the character after the last character of the compiled regular expression. This is useful to programs which have memory allocation to manage.
ERROR(val)	This is the abnormal return from the <i>compile</i> routine. The argument <i>val</i> is an error number (see table below for meanings). This call should never return.

## REGEXP(5)

ERROR	MEANING
11	Range endpoint too large.
16	Bad number.
25	"\digit" out of range.
36	Illegal or missing delimiter.
41	No remembered search string.
42	\(\) imbalance.
43	Too many \(.
44	More than 2 numbers given in \(\).
45	) expected after \.
46	First number exceeds second in \(\).
49	imbalance.
50	Regular expression overflow.

The syntax of the *compile* routine is as follows:

```
compile(instring, expbuf, endbuf, eof)
```

The first parameter *instring* is never used explicitly by the *compile* routine but is useful for programs that pass down different pointers to input characters. It is sometimes used in the INIT declaration (see below). Programs which call functions to input characters or have characters in an external array can pass down a value of (char \*) 0 for this parameter.

The next parameter *expbuf* is a character pointer. It points to the place where the compiled regular expression will be placed.

The parameter *endbuf* is one more than the highest address where the compiled regular expression may be placed. If the compiled expression cannot fit in (*endbuf-expbuf*) bytes, a call to ERROR(50) is made.

The parameter *eof* is the character which marks the end of the regular expression. For example, in *ed*(1), this character is usually a */*.

Each program that includes this file must have a *#define* statement for INIT. This definition will be placed right after the declaration for the function *compile* and the opening curly brace (*{*). It is used for dependent declarations and initializations. Most often it is used to set a register variable to point the beginning of the regular expression so that this register variable can be used in the declarations for GETC(), PEEKC() and UNGETC(). Otherwise it can be used to declare external variables that might be used by GETC(), PEEKC() and UNGETC(). See the example below of the declarations taken from *grep*(1).

There are other functions in this file which perform actual regular expression matching, one of which is the function *step*. The call to *step* is as follows:

```
step(string, expbuf)
```

The first parameter to *step* is a pointer to a string of characters to be checked for a match. This string should be null terminated.

The second parameter *expbuf* is the compiled regular expression which was obtained by a call of the function *compile*.

The function *step* returns non-zero if the given string matches the regular expression, and zero if the expressions do not match. If there is a match, two external character pointers are set as a side effect to the call to *step*. The variable set in *step* is *loc1*. This is a pointer to the first character that matched the regular expression. The variable *loc2*, which is set by the function *advance*, points to the character after the last character that matches the regular expression. Thus if the regular expression matches the entire line, *loc1* will point to the first character of *string* and *loc2* will point to the null at the end of *string*.

*Step* uses the external variable *circf* which is set by *compile* if the regular expression begins with *^*. If this is set then *step* will try to match the regular expression to the beginning of the string only. If more than one regular expression is to be compiled before the first is executed the value of *circf* should be saved for each compiled expression and *circf* should be set to that saved value before each call to *step*.

The function *advance* is called from *step* with the same arguments as *step*. The purpose of *step* is to step through the *string* argument and call *advance* until *advance* returns non-zero indicating a match or until the end of *string* is reached. If one wants to constrain *string* to the beginning of the line in all cases, *step* need not be called; simply call *advance*.

When *advance* encounters a *\** or *\{^}* sequence in the regular expression, it will advance its pointer to the string to be matched as far as possible and will recursively call itself trying to match the rest of the string to the rest of the regular expression. As long as there is no match, *advance* will back up along the string until it finds a match or reaches the point in the string that initially matched the *\** or *\{^}*. It is sometimes desirable to stop this backing up before the initial point in the string is reached. If the external character pointer *locs* is equal to the point in the string at sometime during the backing up process, *advance* will break out of the loop that backs up and will return zero. This is used by *ed(1)* and *sed(1)* for substitutions done globally (not just the first occurrence, but the whole line) so, for example, expressions like *s/y\*/g* do not loop forever.

The additional external variables *sed* and *nbra* are used for special purposes.

#### EXAMPLES

The following is an example of how the regular expression macros and calls look from *grep(1)*:

```
#define INIT    register char *sp = instring;
#define GETC() (*sp++)
#define PEEKC()  (*sp)
#define UNGETC(c) (--sp)
#define RETURN(c)  return;
#define ERROR(c)  regerr()

#include <regexp.h>
...
                (void) compile(*argv, expbuf, &expbuf[ESIZE], \0);
...
                if (step(linebuf, expbuf))
                    succeed();
```

#### FILES

/usr/include/regexp.h

#### SEE ALSO

*bs(1)*, *ed(1)*, *expr(1)*, *grep(1)*, *sed(1)* in the *Runtime System manual*.

#### BUGS

The handling of *circf* is kludgy.

The actual code is probably easier to understand than this manual page.

# STAT(5)

## NAME

stat -- data returned by stat system call

## SYNOPSIS

```
#include <sys/types.h>
#include <sys/stat.h>
```

## DESCRIPTION

The system calls *stat* and *fstat* return data whose structure is defined by this include file. The encoding of the field *st\_mode* is defined in this file also.

```
/*
 * Structure of the result of stat
 */
```

```
struct stat
{
    dev_t    st_dev;
    ino_t    st_ino;
    ushort  st_mode;
    short    st_nlink;
    ushort  st_uid;
    ushort  st_gid;
    dev_t    st_rdev;
    off_t    st_size;
    time_t   st_atime;
    time_t   st_mtime;
    time_t   st_ctime;
};
```

```
#define S_IFMT    0170000 /* type of file */
#define S_IFDIR   0040000 /* directory */
#define S_IFCHR   0020000 /* character special */
#define S_IFBLK   0060000 /* block special */
#define S_IFREG   0100000 /* regular */
#define S_IFIFO   0010000 /* fifo */
#define S_ISUID   04000 /* set user id on execution */
#define S_ISGID   02000 /* set group id on execution */
#define S_ISVTX   01000 /* save swapped text even after use */
#define S_IRREAD  00400 /* read permission, owner */
#define S_IWRITE  00200 /* write permission, owner */
#define S_IXEXEC  00100 /* execute/search permission, owner */
```

## FILES

```
/usr/include/sys/types.h
/usr/include/sys/stat.h
```

## SEE ALSO

stat(2), types(5).

## NAME

term — conventional names for terminals

## DESCRIPTION

These names are used by certain commands [e.g., *tabs(1)*, *man(1)*] and are maintained as part of the shell environment [see *sh(1)*, *profile(4)*, and *envron(5)*] in the variable *\$TERM*:

1520	Datamedia 1520
1620	DIABLO 1620 and others using the HyType II printer
1620-12	same, in 12-pitch mode
2621	Hewlett-Packard HP2621 series
2631	Hewlett-Packard 2631 line printer
2631-c	Hewlett-Packard 2631 line printer - compressed mode
2631-e	Hewlett-Packard 2631 line printer - expanded mode
2640	Hewlett-Packard HP2640 series
2645	Hewlett-Packard HP264n series (other than the 2640 series)
300	DASI/DTC/GSI 300 and others using the HyType I printer
300-12	same, in 12-pitch mode
300s	DASI/DTC/GSI 300s
382	DTC 382
300s-12	same, in 12-pitch mode
3045	Datamedia 3045
33	TELETYPE® Model 33 KSR
37	TELETYPE Model 37 KSR
40-2	TELETYPE Model 40/2
40-4	TELETYPE Model 40/4
4540	TELETYPE Model 4540
3270	IBM Model 3270
4000a	Trendata 4000a
4014	TEKTRONIX 4014
43	TELETYPE Model 43 KSR
450	DASI 450 (same as DIABLO 1620)
450-12	same, in 12-pitch mode
735	Texas Instruments TI735 and TI725
745	Texas Instruments TI745
dumb	generic name for terminals that lack reverse line-feed and other special escape sequences
sync	generic name for synchronous TELETYPE 4540-compatible terminals
hp	Hewlett-Packard (same as 2645)
lp	generic name for a line printer
tn1200	User Electric TermiNet 1200
tn300	User Electric TermiNet 300

Up to 8 characters, chosen from [-a-z0-9], make up a basic terminal name. Terminal sub-models and operational modes are distinguished by suffixes beginning with a -. Names should generally be based on original vendors, rather than local distributors. A terminal acquired from one vendor should not have more than one distinct basic name.

Commands whose behavior depends on the type of terminal should accept arguments of the form *-Term* where *term* is one of the names given above; if no such argument is present, such commands should obtain the terminal type from the environment variable *\$TERM*, which, in turn, should contain *term*.

## TERM(5)

### SEE ALSO

`profile(4)`, `environ(5)`.

`man(1)`, `tpplot(1G)`, `sh(1)`, `stty(1)`, `tabs(1)` in the Runtime System manual.

`mm(1)`, `nroff(1)` in the "User Reference Manual" chapter of the Text Preparation System manual.

### BUGS

This is a small candle trying to illuminate a large, dark problem. Programs that ought to adhere to this nomenclature do so somewhat fitfully.

## NAME

types — primitive system data types

## SYNOPSIS

#include &lt;sys/types.h&gt;

## DESCRIPTION

The data types defined in the include file are used in UNIX system code; some data of these types are accessible to user code:

```

typedef struct { int r[1]; } *      physadr;
typedef long      daddr_t;
typedef char *    caddr_t;
typedef unsigned int  uint;
typedef unsigned short ushort;
typedef ushort    ino_t;
typedef short     cnt_t;
typedef long      time_t;
typedef int       label_t[10];
typedef short     dev_t;
typedef long      off_t;
typedef long      paddr_t;
typedef long      key_t;

```

The form *daddr\_t* is used for disk addresses except in an i-node on disk, see *fs(4)*. Times are encoded in seconds since 00:00:00 GMT, January 1, 1970. The major and minor parts of a device code specify kind and unit number of a device and are installation-dependent. Offsets are measured in bytes from the beginning of a file. The *label\_t* variables are used to save the processor state while another process is running.

## SEE ALSO

*fs(4)*.

## VALUES(5)

### NAME

values — machine-dependent values

### SYNOPSIS

```
#include <values.h>
```

### DESCRIPTION

This file contains a set of manifest constants, conditionally defined for particular processor architectures.

The model assumed for integers is binary representation (one's or two's complement), where the sign is represented by the value of the high-order bit.

<b>BITS</b> ( <i>type</i> )	The number of bits in a specified type (e.g., int).
<b>HIBITS</b>	The value of a short integer with only the high-order bit set (in most implementations, 0x8000).
<b>HIBITL</b>	The value of a long integer with only the high-order bit set (in most implementations, 0x80000000).
<b>HIBITI</b>	The value of a regular integer with only the high-order bit set (usually the same as HIBITS or HIBITL).
<b>MAXSHORT</b>	The maximum value of a signed short integer (in most implementations, 0x7FFF $\equiv$ 32767).
<b>MAXLONG</b>	The maximum value of a signed long integer (in most implementations, 0x7FFFFFFF $\equiv$ 2147483647).
<b>MAXINT</b>	The maximum value of a signed regular integer (usually the same as MAXSHORT or MAXLONG).
<b>MAXFLOAT, LN_MAXFLOAT</b>	The maximum value of a single-precision floating-point number, and its natural logarithm.
<b>MAXDOUBLE, LN_MAXDOUBLE</b>	The maximum value of a double-precision floating-point number, and its natural logarithm.
<b>MINFLOAT, LN_MINFLOAT</b>	The minimum positive value of a single-precision floating-point number, and its natural logarithm.
<b>MINDOUBLE, LN_MINDOUBLE</b>	The minimum positive value of a double-precision floating-point number, and its natural logarithm.
<b>FSIGNIF</b>	The number of significant bits in the mantissa of a single-precision floating-point number.
<b>DSIGNIF</b>	The number of significant bits in the mantissa of a double-precision floating-point number.

### FILES

/usr/include/values.h

### SEE ALSO

intro(3), math(5).



**NAME**  
varargs — handle variable argument list

**SYNOPSIS**

```
#include <varargs.h>
va_alist
va_dcl
void va_start(pvar)
va_list pvar;
type va_arg(pvar, type)
va_list pvar;
void va_end(pvar)
va_list pvar;
```

**DESCRIPTION**

This set of macros allows portable procedures that accept variable argument lists to be written. Routines that have variable argument lists [such as *printf*(3S)] but do not use *varargs* are inherently nonportable, as different machines use different argument-passing conventions.

*va\_alist* is used as the parameter list in a function header.

*va\_dcl* is a declaration for *va\_alist*. No semicolon should follow *va\_dcl*.

*va\_list* is a type defined for the variable used to traverse the list.

*va\_start* is called to initialize *pvar* to the beginning of the list.

*va\_arg* will return the next argument in the list pointed to by *pvar*. *Type* is the type the argument is expected to be. Different types can be mixed, but it is up to the routine to know what type of argument is expected, as it cannot be determined at runtime.

*va\_end* is used to clean up.

Multiple traversals, each bracketed by *va\_start* ... *va\_end*, are possible.

**EXAMPLE**

This example is a possible implementation of *execl*.

```
#include <varargs.h>
#define MAXARGS 100

/* execl is called by
   execl(file, arg1, arg2, ..., (char *)0);
*/
execl(va_alist)
va_dcl
{
    va_list ap;
    char *file;
    char *args[MAXARGS];
    int argno = 0;

    va_start(ap);
    file = va_arg(ap, char *);
    while ((argno++) != va_arg(ap, char *)) != (char *)0
        ;
    va_end(ap);
    return execv(file, args);
}
```

## VARARGS(5)

### SEE ALSO

`exec(2)`, `printf(3S)`.

### BUGS

It is up to the calling routine to specify how many arguments there are, since it is not always possible to determine this from the stack frame. For example, `execl` is passed a zero pointer to signal the end of the list. `Printf` can tell how many arguments are there by the format.

It is non-portable to specify a second argument of *char*, *short*, or *float* to `va_arg`, since arguments seen by the called function are not *char*, *short*, or *float*. C converts *char* and *short* arguments to *int* and converts *float* arguments to *double* before passing them to a function.

**NOTES**



## **NOTES**

**NOTES**



**NOTES**

**NOTES**

U

O

C

**NOTES**

( )

( )

( )