

UNIPLUS⁺ SYSTEM V

Users's Manual

Sections 2 – 6

UniSoft

INTRODUCTION

This manual describes the features of System V UniPlus⁺, a UNIX operating system. All commands, features, and facilities described in this manual are available on UniPlus⁺.

This manual is divided into two volumes containing a total of six sections, some containing subsections:

1. Commands and Application Programs:
 1. General-Purpose Commands.
 - 1C. Communications Commands.
 - 1G. Graphics Commands.
2. System Calls.
3. Subroutines:
 - 3C. C and Assembler Library Routines.
 - 3M. Mathematical Library Routines.
 - 3S. Standard I/O Library Routines.
 - 3X. Miscellaneous Routines.
4. File Formats.
5. Miscellaneous Facilities.
6. Games.

Section 1 (*Commands and Application Programs*) describes programs intended to be invoked directly by the user or by command language procedures, as opposed to subroutines, which are intended to be called by the user's programs. Commands generally reside in the directory `/bin` (for **binary** programs). Some programs also reside in `/usr/bin`, to save space in `/bin`. These directories are searched automatically by the command interpreter called the *shell*. Sub-class 1C contains communication programs such as *cu*, *send*, *uucp*, etc.

Section 2 (*System Calls*) describes the entries into the UNIX kernel, including the C language interface.

Section 3 (*Subroutines*) describes the available subroutines. Their binary versions reside in various system libraries in the directories `/lib` and `/usr/lib`. See *intro(3)* for descriptions of these libraries and the files in which they are stored.

Section 4 (*File Formats*) documents the structure of particular kinds of files; for example, the format of the output of the link editor is given in *a.out(4)*. Excluded are files used by only one command (for example, the assembler's intermediate files). In general, the C language `struct` declarations corresponding to these formats can be found in the directories `/usr/include` and `/usr/include/sys`.

Section 5 (*Miscellaneous Facilities*) includes descriptions of character sets, macro packages and other system features.

Section 6 (*Games*) describes the games and educational programs that, as a rule, reside in the directory `/usr/games`.

Each section consists of a number of independent entries of a page or so each. The name of the entry appears in the upper corners of its pages. Entries within each

Copyright © 1983 UniSoft Corporation.

Portions of this material have been previously copyrighted by:

Bell Telephone Laboratories, Incorporated, 1980

Western Electric Company, Incorporated, 1983

Regents of the University of California

Holders of a UNIX and UniPlus⁺ software license are permitted to copy this document, or any portion of it, as necessary for licensed use of the software, provided this copyright notice and statement of permission are included.

UNIX is a Trademark of Bell Telephone Laboratories, Inc.

UniPlus⁺ is a Trademark of UniSoft Corporation of Berkeley.

section are alphabetized, with the exception of the introductory entry that begins each section. The page numbers of each entry start at 1. The version date of the entry appears in the lower left corner of each page. Some entries may describe several routines, commands, etc. In such cases, the entry appears only once, alphabetized under its "major" name.

All entries are based on a common format, not all of whose parts always appear:

- The **NAME** part gives the name(s) of the entry and briefly states its purpose.
- The **SYNOPSIS** part summarizes the use of the program being described. A few conventions are used, particularly in Section 1 (*Commands*):

- Boldface** strings are literals and are to be typed just as they appear.
- Italic* strings usually represent substitutable argument prototypes and program names found elsewhere in the manual.
- Square brackets **[]** around an argument prototype indicate that the argument is optional. When an argument prototype is given as "name" or "file", it always refers to a *file* name.
- Ellipses **...** are used to show that the previous argument prototype may be repeated.
- A final convention is used by the commands themselves. An argument beginning with a minus **-**, plus **+**, or equal sign **=** is often taken to be some sort of flag argument, even if it appears in a position where a file name could appear. Therefore, it is unwise to have files whose names begin with **-**, **+**, or **=**.

- The **DESCRIPTION** part discusses the subject at hand.
- The **EXAMPLE** part gives example(s) of usage, where appropriate.
- The **FILES** part gives the file names that are built into the program.
- The **SEE ALSO** part gives pointers to related information.
- The **DIAGNOSTICS** part discusses the diagnostic indications that may be produced. Messages that are intended to be self-explanatory are not listed.
- The **WARNINGS** part points out potential pitfalls.
- The **BUGS** part gives known bugs and sometimes deficiencies. Occasionally, the suggested fix is also described.

At the front of each volume there is a table of contents and a permuted index. The permuted index is a computer-generated index that uses the information in the **NAME** part of each entry in the User's and Administrator's Manuals. The permuted index contains three columns. The center column is an alphabetic list of keywords as they appear in the **NAME** part of the entries. The last column is the entry that the keyword in the center column refers to. This entry is followed by the appropriate section number in parentheses. The first column contains the remaining information from the **NAME** part that either precedes or follows the keyword.

For example, to look for a text editor, scan the center column for the word "editor". There are several index lines containing an "editor" reference, i.e.:

```
ed, red: text editor. . . . . ed(1)
files. ld: link editor for common object . . . . . ld(1)
```

You can then turn to the entries listed in the last column, *ed(1)* and *ld(1)*, to find information on the editor.

On most systems, all user manual entries are available on-line via the command, q.v.

TABLE OF CONTENTS

2. System Calls

intro	introduction to system calls and error numbers
accept	accept a connection on a socket
access	determine accessibility of a file
acct	enable or disable process accounting
alarm	set a process's 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
connect	initiate a connection on a socket
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
gethostname	get name of current host
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
lockf	provide exclusive file regions for reading or writing
lseek	move read/write file pointer
mknod	make a directory, or a special 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
phys	allow a process to access physical addresses
pipe	create an interprocess channel
plock	lock process, text, or data in memory
profil	execution time profile
ptrace	process trace
read	read from file
reboot	reboot the system
receive	receive message from a socket
select	synchronous i/o multiplexing
semctl	semaphore control operations
semget	get set of semaphores
semop	semaphore operations
send	send message from a socket
sethostname	set name of host cpu
setpgrp	set process group ID
setuid	set user and group IDs
shmctl	shared memory control operations
shmget	get shared memory segment
shmop	shared memory operations

signal	specify what to do upon receipt of a signal
socket	create an endpoint for communication
socketaddr	return address associated with a socket
stat	get file status
stime	set time
sync	update super-block
time	get time
times	get process and child process times
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
uvar	returns system-specific configuration information
wait	wait for child process to stop or terminate
write	write on a file

3. Subroutines

intro	introduction to subroutines and libraries
a64l	convert between long integer and base-64 ASCII string
abort	generate an IOT fault
abs	return integer absolute value
assert	verify program assertion
atof	convert ASCII string to floating-point number
bessel	Bessel functions
blt	block transfer data
bsearch	binary search
clock	report CPU time used
conv	translate characters
crypt	generate DES encryption
ctermid	generate file name for terminal
ctime	convert date and time to string
ctype	classify characters
cuserid	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
erf	error function and complementary error function
exp	exponential, logarithm, power, square root functions
fclose	close or flush a stream
ferror	stream status inquiries
floor	floor, ceiling, remainder, absolute value functions
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
gamma	log gamma function
getc	get character or word from 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
hypot	Euclidean distance function
l3tol	convert between 3-byte integers and long integers
logname	return login name of user
lsearch	linear search and update
malloc	main memory allocator
matherr	error-handling function
memory	memory operations
mktemp	make a unique file name
monitor	prepare execution profile
nlist	get entries from name list
perror	system error messages
plot	graphics interface subroutines
popen	initiate pipe to/from a process
printf	print formatted output
putc	put character or word on a stream
putpwent	write password file entry
puts	put a string on a stream
qsort	quicker sort
rand	simple random-number generator
regcmp	compile and execute regular expression
rhost	look up internet hosts by name or address
scanf	convert formatted input
setbuf	assign buffering to a stream
setjmp	non-local goto
sinh	hyperbolic functions
sleep	suspend execution for interval
sputl	access long numeric data in a machine independent fashion.
signal	software signals
stdio	standard buffered input/output package
stdipc	standard interprocess communication package
string	string operations
strtol	convert string to integer
swab	swap bytes
system	issue a shell command
termcap	terminal independent operation routines
tmpfile	create a temporary file
tmpnam	create a name for a temporary file
trig	trigonometric functions
tsearch	manage binary search trees
ttyname	find name of a terminal
ttyslot	find the slot in the utmp file of the current user
ungetc	push character back into input stream

4. File Formats

intro	introduction to file formats
a.out	assembler and link editor output
acct	per-process accounting file format
altblk	alternate block information for bad block handling
ar	archive (library) file format

checklist list of file systems processed by fsck
 core format of core image file
 cpio format of cpio archive
 dir format of directories
 environ user environment
 errfile error-log file format
 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 inode
 issue issue identification file
 master master device information table
 mnttab mounted file system table
 passwd password file
 plot graphics interface
 pnch file format for card images
 profile setting up an environment at login time
 sccsfile format of SCCS file
 tp magnetic tape format
 ttytype data base of terminal types by port
 utmp utmp and wtmp entry formats

5. Miscellaneous Facilities

intro introduction to miscellany
 ascii map of ASCII character set
 environ user environment
 eqnchar special character definitions for eqn and neqn
 fcntl file control options
 greek graphics for the extended TTY-37 type-box
 inet Internet protocol family
 ip Internet Protocol
 lo software loopback interface
 man macros for formatting entries in this manual
 mm the MM macro package for formatting documents
 mosd the OSDD adapter macro package for formatting documents
 mptx the macro package for formatting a permuted index
 mv a troff macro package for typesetting view graphs and slides
 net introduction to networking facilities
 regexp regular expression compile and match routines
 stat data returned by stat system call
 tcp Internet Transmission Control Protocol
 term conventional names for terminals
 termcap terminal capability data base
 types primitive system data types
 udp Internet User Datagram Protocol

6. Games

intro introduction to games
 adventure an exploration game
 aliens The alien invaders attack the earth
 arithmetic provide drill in number facts
 autorobots Escape from the automatic robots
 back the game of backgammon

bcd convert to antique media
 bj the game of black jack
 chase Try to escape the killer robots
 craps the game of craps
 cribbage the card game cribbage
 fish play "Go Fish"
 fortune print a random, hopefully interesting, adage
 hangman guess the word
 life play the game of life
 maze generate a maze
 moo guessing game
 number convert Arabic numerals to English
 quiz test your knowledge
 rain animated raindrops display
 robots Escape from the robots
 trek trekkie game
 ttt tic-tac-toe
 twinkle twinkle stars on the screen
 worm Play the growing worm game
 worms animate worms on a display terminal
 wump the game of hunt-the-wumpus

PERMUTED INDEX

/functions of HP 2640 and 2621-series terminals. hp.1
 handle special functions of HP 2640 and 2621-series/ hp: hp.1
 archiver. hpio: HP 2645A terminal tape file hpio.1
 300, 300s: handle special 300.1
 functions of DASI 300 and/ 300 and 300s terminals. 300.1
 /special functions of DASI 300s: handle special functions 300.1
 of DASI 300 and 300s/ 300, 300s terminals. /special 300.1
 functions of DASI 300 and 3-byte integers and long/ l3tol.3c
 l3tol, ltol3: convert between 3-way differential file diff3.1
 comparison. diff3: 4014: paginator for the 4014.1
 Tektronix 4014 terminal. 4014 terminal. 4014: 4014.1
 paginator for the Tektronix 450: handle special functions 450.1
 of the DASI 450 terminal. 450 terminal. 450: handle 450.1
 special functions of the DASI a64l, l64a: convert between a64l.3c
 long integer and base-64/ abort: generate an IOT fault. abort.3c
 value. abs: return integer absolute abs.3c
 abs: return integer absolute value. abs.3c
 /floor, ceiling, remainder, absolute value functions. floor.3m
 socket. accept: accept a connection on a accept.2
 a socket. accept: accept a connection on accept.2
 LP requests. accept, reject: allow/prevent accept.1m
 utime: set file access and modification times. utime.2
 of a file. touch: update access and modification times touch.1
 accessibility of a file. access: determine access.2
 machine/ sputl, sgetl: access long numeric data in a sputl.3x
 phys: allow a process to access physical addresses. phys.2
 sadb: disk access profiler. sadp.1
 copy file systems for optimal access time. dcopy: dcopy.1m
 /setutent, endutent, utmpname: access utmp file entry. getut.3c
 access: determine accessibility of a file. access.2
 enable or disable process accounting. acct: acct.2
 acctcon2: connect-time accounting. acctcon1, acctcon.1m
 acctprc1, acctprc2: process accounting. acctprc.1m
 turnacct: shell procedures for acctsh.1m
 runacct: run daily accounting. /startup, runacct.1m
 /accounton, acctwtmp: overview of accounting. acct.1m
 accounting and miscellaneous/ accounting commands. /of acct.1m
 accounting file format. acct.4
 acct: per-process accounting file(s). acctcom: acctcom.1
 search and print process accounting files. acctmerg.1m
 acctmerg: merge or add total accounting records. /command acctcms.1m
 summary from per-process accounting records. fwtmp, fwtmp.1m
 wtmpfix: manipulate connect acct: enable or disable acct.2
 process accounting. acct: per-process accounting acct.4
 file format. acctcms: command summary from acctcms.1m
 per-process accounting/ acctcms: search and print acctcom.1
 process accounting file(s). acctcom: search and print acctcon.1m
 connect-time accounting. acctcon1, acctcon2: acctcon.1m
 accounting. acctcon1, acctcon2: connect-time acct.1m
 acctwtmp: overview of/ acctdisk, acctdusg, accton, acct.1m
 overview of/ acctdisk, acctdusg, accton, acctwtmp: acctmerg.1m
 acctmerg: merge or add total accounting files. acct.1m
 acctdisk, acctdusg, accton, acctwtmp: overview of/ acctprc.1m
 accounting. acctprc1, acctprc2: process acctprc.1m
 acctprc1, acctprc2: process accounting. acct.1m
 acctwtmp: overview of/ acctwtmp: overview of/ trig.3m
 acctdisk, acctdusg, accton, acos, atan, atan2:/ killall.1m
 sin, cos, tan, asin, active processes. sact.1
 killall: kill all activity. sact: print timex.1
 current SCCS file editing activity. /time a command; sag.1g
 report process data and system activity graph. sar.1m
 sag: system activity report package. sar.1m
 sa1, sa2, sadc: system

random, hopefully interesting, formatting/ mosd: the OSDD
 acctmerg: merge or up internet hosts by name or socket. socketaddr: return a process to access physical SCCS files.
 admin: create and game. alarm: set a process's clock. delivermail. aliases: earth. aliens: The attack the earth. change data segment space realloc, calloc: main memory physical addresses. phys: accept, reject: information for bad block/ for bad block/ altblk: sort: terminal. worms: rain: bcd: convert to editor output. introduction to commands and maintenance commands and maintainer. format. number: convert delivermail: deliver mail to language. bc: cpio: format of cpio tp: manipulate tape maintainer. ar: HP 2645A terminal tape file tar: tape file cpio: copy file command. xargs: construct getopt: get option letter from echo: echo expr: evaluate bc: arbitrary-precision number facts. expr: evaluate arguments characters. asa: interpret control characters. ascii: map of /translates object files into set. long integer and base-64 number. atof: convert and/ ctime, localtime, gmtime, trigonometric/ sin, cos, tan, help: as: output. a.out: assertion. assert: verify program
 activity reporter. sar.1
 adage. fortune: print a fortune.6
 adapter macro package for mosd.5
 adb: debugger. adb.1
 add total accounting files. acctmerg.1m
 address. rhost, raddr: look rhost.3
 address associated with a socketaddr.2
 addresses. phys: allow phys.2
 admin: create and administer administer SCCS files. admin.1
 admin.1
 adventure: an exploration adventure.6
 alarm clock. alarm.2
 alarm: set a process's alarm alarm.2
 aliases: aliases file for aliases.7
 aliases file for delivermail. aliases.7
 alien invaders attack the aliens.6
 aliens: The alien invaders aliens.6
 allocation. brk, sbrk: brk.2
 allocator. malloc, free, malloc.3c
 allow a process to access phys.2
 allow/prevent LP requests. accept.1m
 altblk: alternate block altblk.4
 alternate block information and/or merge files. sort.1
 animate worms on a display worms.6
 animated raindrops display. rain.6
 antique media. bcd.6
 a.out: assembler and link a.out.4
 application programs. intro: intro.1
 application programs. /system intro.1m
 ar: archive and library ar.1
 ar: archive (library) file ar.4
 Arabic numerals to English. number.6
 arbitrary people. delivermail.8
 arbitrary-precision arithmetic bc.1
 archive. cpio.4
 archive. tp.1
 archive and library ar.1
 archive (library) file format. ar.4
 archiver. hpio: hpio.1
 archiver. tar.1
 archives in and out. cpio.1
 argument list(s) and execute xargs.1
 argument vector. getopt.3c
 arguments. echo.1
 arguments as an expression. expr.1
 arithmetic language. bc.1
 arithmetic: provide drill in arithmetic.6
 as an expression. expr.1
 as: assembler. as.1
 ASA carriage control asa.1
 asa: interpret ASA carriage asa.1
 ASCII character set. ascii.5
 ASCII formats suitable for/ hex.1
 ascii: map of ASCII character ascii.5
 ASCII string. /convert between a64l.3c
 ASCII string to floating-point atof.3c
 asctime, tzset: convert date ctime.3c
 asin, acos, atan, atan2: trig.3m
 ask for help. help.1
 assembler. as.1
 assembler and link editor a.out.4
 assert: verify program assert.3x
 assertion. assert.3x

setbuf: assign buffering to a stream. setbuf.3s
 *associated with a socket. socketaddr.2
 atan, atan2: trigonometric/ trig.3m
 atan2: trigonometric/ sin, trig.3m
 cos, tan, asin, acos, atan,
 floating-point number.
 atof: convert ASCII string to atof.3c
 atoi: convert string to strtol.3c
 atol, atoi: convert string to strtol.3c
 integer. strtol, atol,
 integer. strtol,
 aliens: The alien invaders
 autorobots: Escape from the autorobots.6
 automatic robots. autorobots.6
 automatic robots. autorobots.6
 wait: await completion of process. wait.1
 processing language.
 ungetc: push character
 back: the game of
 backgammon. back.6
 backup. filesave, tapesave: filesave.1m
 backup. fnc.1m
 backup tape. freq.1m
 bad block handling. /alternate altblk.4
 bad block information. badblk.1m
 badblk: program to set or badblk.1m
 banner: make posters. banner.1
 banner7: print large banner7.1
 printer. banner7.1
 terminal capability data
 port. ttytype: data termcap.5
 base of terminal types by ttytype.4
 base-64 ASCII string. /convert a64l.3c
 based on ex. /screen oriented vi.1
 basename, dirname: deliver basename.1
 bc: arbitrary-precision bc.1
 bcd: convert to antique media. bcd.6
 bcheckrc, rc, powerfail: brc.1m
 bcopy: interactive block copy. bcopy.1m
 bdiff: big diff. bdiff.1
 beautifier. cb.1
 (Berkeley version). ls7: ls7.1
 Bessel functions. bessel.3m
 bfs: big file scanner. bfs.1
 strings in an object, or other
 fread, fwrite: strings.1
 bsearch: fread.3s
 binary search. bsearch.3c
 binary file. /the printable tsearch.3c
 binary input/output.
 tdelete, twalk: manage
 remove symbols and relocation
 reset: reset the teletype
 reset.
 bj: the game of black jack. bj.6
 black jack. bj.6
 block. sync.1
 bcopy: interactive bcopy.1m
 block copy. sum.1
 block count of a file.
 block handling. /alternate altblk.4
 block information for bad badblk.1m
 program to set or update bad
 block/ altblk: alternate altblk.4
 block information for bad altblk.4
 blt, blt512: blt.3
 block transfer data. blt.3
 blocks. df.1m
 sum7: sum and count sum7.1
 data. blt.3
 blt, blt.3
 blt512: block transfer data. blt.3
 /etc/hosts: host table for hosts.7
 netmail: the netmail.8
 bnet network mail system. boot.8
 boot: startup procedures. boot.8
 brc, bcheckrc, rc, powerfail: brc.1m
 brk, sbrk: change data segment brk.2
 bs: a compiler/interpreter for bs.1
 bsearch: binary search. bsearch.3c

stdio: standard buffered input/output package. stdio.3s
 setbuf: assign buffering to a stream. setbuf.3s
 mknod: build special file. mknod.1m
 swab: swap bytes. swab.3c
 cc: C compiler. cc.1
 cflow: generate C flow graph. cflow.1
 cpp: the C language preprocessor. cpp.1
 maintain a tags file for a C program. ctags: ctags.1
 cb: C program beautifier. cb.1
 lint: a C program checker. lint.1
 cxref: generate C program cross reference. cxref.1
 message file by massaging C source. /create an error mkstr.1
 dc: desk cal: print calendar. cal.1
 calculator. dc.1
 calendar. cal.1
 calendar: reminder service. calendar.1
 call. stat: stat.5
 data returned by stat system cu: call another UNIX System. cu.1c
 malloc, free, realloc, link and unlink system
 intro: introduction to system calls and error numbers. intro.2
 to an LP line printer. lp, cancel: send/cancel requests lp.1
 termcap: terminal capability data base. termcap.5
 cribbage: the card game cribbage. cribbage.6
 pnch: file format for card images. pnch.4
 asa: interpret ASA carriage control characters. asa.1
 files. cat: concatenate and print cat.1
 cb: C program beautifier. cb.1
 cc: C compiler. cc.1
 cd: change working directory. cd.1
 commentary of an SCCS delta. cdc: change the delta cdc.1
 ceiling, remainder,/ floor, ceil, fmod, fabs: floor, /ceil, fmod, fabs: floor, ceiling, remainder, absolute/ floor.3m
 delta: make a delta cflow: generate C flow graph. cflow.1
 pipe: create an interprocess (change) to an SCCS file. delta.1
 stream. ungetc: push channel. pipe.2
 and neqn. eqnchar: special character back into input ungetc.3s
 file. freq: report on character definitions for eqn eqnchar.5
 user. cuserid: get character frequencies in a freq.1
 /getchar, fgetc, getw: get character login name of the cuserid.3s
 /putchar, fputc, putw: put character or word from stream. getc.3s
 ascii: map of ASCII character or word on a stream. putc.3s
 interpret ASA carriage control character set. ascii.5
 _tolower, toascii: translate characters. asa: asa.1
 iscntrl, isascii: classify characters. /_toupper, conv.3c
 tr: translate characters. /isprint, isgraph, ctype.3c
 given/ sumdir: sum and count characters. tr.1
 lastlogin, monacct, nulladm,/ characters in the files in the sumdir.1
 killer robots. chargefee, ckpacct, dodisk, acctsh.1m
 directory. chase: Try to escape the chase.6
 /dfsck: file system consistency chdir: change working chdir.2
 constant-width text for/ cw, check and interactive repair. fsck.1m
 text for nroff or/ eqn, neqn, checker. lint.1
 lint: a C program checkcw: prepare cw.1
 grpck: password/group file checkeq: format mathematical eqn.1
 copy file systems with label checker. lint.1
 systems processed by fsck. pwck: pwck.1m
 formatted with the/ mm, osdd, checking. volcopy, labelit: volcopy.1m
 file. sum: print checklist: list of file checklist.4
 vchk: version checkmm: print/check documents mm.1
 chown, checksum and block count of a sum.1
 times: get process and checkup. vchk.1m
 terminate. wait: wait for chgrp: change owner or group. chown.1
 child process times. times.2
 child process to stop or wait.2

chmod: change mode. chmod.1
 chmod: change mode of file. chmod.2
 of a file. chown: change owner and group chown.2
 group. chown, chgrp: change owner or chown.1
 for a command. chroot: change root directory. chroot.2
 monacct, nulladm,/ chargefee, chroot: change root directory chroot.1m
 isgraph, iscntrl, isascii: ckpacct, dodisk, lastlogin, acctsh.1m
 uuclean: uucp spool directory classify characters. /isprint, ctype.3c
 clean-up. uuclean.1m
 clear: clear terminal screen. clear.1
 clri: clear i-node. clri.1m
 clear: clear terminal screen. clear.1
 status/ ferror, feof, clearerr, fileno: stream ferror.3s
 (command interpreter) with C-like syntax. csh: a shell csh.1
 alarm: set a process's alarm clock. alarm.2
 cron: clock daemon. cron.1m
 clock: report CPU time used. clock.3c
 close: close a file descriptor. close.2
 descriptor. close: close a file close.2
 fclose, fflush: close or flush a stream. fclose.3s
 clri: clear i-node. clri.1m
 cmp: compare two files. cmp.1
 line-feeds. col: filter reverse col.1
 comb: combine SCCS deltas. comb.1
 common to two sorted files. combine SCCS deltas. comb.1
 change root directory for a comm: select or reject lines comm.1
 system: issue a shell command. chroot: chroot.1m
 test: condition evaluation system. system.3s
 time: time a command. test.1
 argument list(s) and execute time: time a command. time.1
 nice: run a command. xargs: construct xargs.1
 env: set environment for nice: run a nice.1
 uux: unix to unix command execution. env.1
 (sh/ nohup: run a command execution. uux.1c
 C-like syntax. csh: a shell command immune to hangups nohup.1
 getopt: parse (command interpreter) with csh.1
 /shell, the standard/restricted command options. getopt.1
 and system/ timex: time a command programming language. sh.1
 per-process/ acctcms: command; report process data timex.1
 and miscellaneous accounting acctcms: command summary from acctcms.1m
 install: install commands. /of accounting acct.1m
 intro: introduction to commands and application/ intro.1m
 /to system maintenance intro: introduction to intro.1m
 at: execute commands at a later time. at.1
 cdc: change the delta. commentary of an SCCS delta. cdc.1
 comm: select or reject lines common to two sorted files. comm.1
 socket: create an endpoint for communication. socket.2
 ipc: report inter-process communication facilities/ ipc.1
 stdipc: standard interprocess communication package. stdipc.3c
 diff: differential file comparator. diff.1
 cmp: compare two files. cmp.1
 SCCS file. sccsdiff: compare two versions of an sccsdiff.1
 diff3: 3-way differential file comparison. diff3.1
 dircmp: directory comparison. dircmp.1
 regcmp: regular expression compile. regcmp.1
 expression. regcmp, regex: compile and execute regular regcmp.3x
 regex: regular expression compile and match routines. regexp.5
 cc: C compiler. cc.1
 fortran: FORTRAN compiler. fortran.1
 yacc: yet another compiler-compiler. yacc.1
 modest-sized programs. bs: a compiler/interpreter for bs.1
 erf, erfc: error function and complementary error function. erf.3m
 wait: await completion of process. wait.1

pack, pcat, unpack: compress and expand files. pack.1
 cat: concatenate and print files. cat.1
 test: condition evaluation command. test.1
 uvar: returns system-specific configuration information. uvar.2
 system. lpadmin: configure the LP spooling lpadmin.1m
 fwtmp, wtmpfix: manipulate connect accounting records. fwtmp.1m
 on a socket. connect: initiate a connection connect.2
 an out-going terminal line connection. dial: establish dial.3c
 accept: accept a connection on a socket. accept.2
 connect: initiate a connection on a socket. connect.2
 acctcon1, acctcon2: connect-time accounting. acctcon.1m
 fsck, dfsc: file system consistency check and/ fsck.1m
 cw, checkcw: prepare constant-width text for troff. cw.1
 mkfs: construct a file system. mkfs.1m
 mkfs512: construct a file system. mkfs512.1m
 execute command. xargs: construct argument list(s) and xargs.1
 nroff/troff, tbl, and eqn constructs. deroff: remove deroff.1
 ls: list contents of directories. ls.1
 (Berkeley version). ls7: list contents of directory ls7.1
 csplit: context split. csplit.1
 fcntl: file control. fcntl.2
 uucp status inquiry and job control. uustat: uustat.1c
 vc: version control. vc.1
 asa: interpret ASA carriage control characters. asa.1
 ioctl: control device. ioctl.2
 init, telinit: process control initialization. init.1m
 msgctl: message control operations. msgctl.2
 semctl: semaphore control operations. semctl.2
 shmctl: shared memory control operations. shmctl.2
 fcntl: file control options. fcntl.5
 tcp: Internet Transmission Control Protocol. tcp.5
 interface. tty: controlling terminal tty.7
 terminals. term: conventional names for term.5
 units: conversion program. units.1
 dd: convert and copy a file. dd.1
 English. number: convert Arabic numerals to number.6
 floating-point number. atof: convert ASCII string to atof.3c
 integers and/ l3tol, l64a: convert between 3-byte l3tol.3c
 and base-64 ASCII/ a64l, l64a: convert between long integer a64l.3c
 /gmtime, asctime, tzset: convert date and time to/ ctime.3c
 to string. ecvt, fcvt, gcvt: convert floating-point number ecvt.3c
 scanf, fscanf, sscanf: convert formatted input. scanf.3s
 strtol, atol, atoi: convert string to integer. strtol.3c
 bcd: convert to antique media. bcd.6
 bcopy: interactive block copy. bcopy.1m
 rcp: remote file copy rcp.1
 uucop, uuname: unix to unix copy. uucp, uucp.1c
 System-to-UNIX System file copy. /uupick: public UNIX uuto.1c
 dd: convert and copy a file. dd.1
 cpio: copy file archives in and out. cpio.1
 access time. dcopy: copy file systems for optimal dcopy.1m
 checking. volcopy: copy file systems with label volcopy.1m
 cp, ln, mv: copy, link or move files. cp.1
 file. core: format of core image core.4
 core image file. core.4
 core memory. mem.7
 atan2: trigonometric/ sin, cos, tan, asin, acos, atan, trig.3m
 functions. sinh, cosh, tanh: hyperbolic sinh.3m
 wc: word count. wc.1
 sum7: sum and count blocks in a file. sum7.1
 in the given/ sumdir: sum and count characters in the files sumdir.1
 sum: print checksum and block count of a file. sum.1
 files. cp, ln, mv: copy, link or move cp.1
 cpio: format of cpio archive. cpio.4

and out. cpio: copy file archives in cpio.1
 cpio: format of cpio archive. cpio.4
 preprocessor. cpp: the C language cpp.1
 cpu. sethostname.2
 sethostname: set name of host clock: report CPU time used. clock.3c
 craps: the game of craps. craps.6
 system crashes. craps: the game of craps. craps.6
 what to do when the system crash. crash: crash.8
 rewrite an existing one. creat: create a new file or creat.2
 file. tmpnam, tempnam: create a name for a temporary tmpnam.3s
 an existing one. creat: create a new file or rewrite creat.2
 fork: create a new process. fork.2
 tmpfile: create a temporary file. tmpfile.3s
 communication. socket: create an endpoint for socket.2
 by messaging C source. mkstr: create an error message file mkstr.1
 channel. pipe: create an interprocess pipe.2
 files. admin: create and administer SCCS admin.1
 umask: set and get file creation mask. umask.2
 cribbage: the card game cribbage. cribbage.6
 cribbage: the card game cribbage.6
 cron: clock daemon. cron.1m
 cxref: generate C program cross reference. cxref.1
 more: file perusal filter for crt viewing. more.1
 generate DES encryption. crypt: encode/decode. crypt.1
 interpreter) with C-like/ crypt, setkey, encrypt: crypt.3c
 csh: a shell (command csh.1
 csplit: context split. csplit.1
 terminal. ct: spawn getty to a remote ct.1c
 for a C program. ctags: maintain a tags file ctags.1
 for terminal. ctermid: generate file name ctermid.3s
 asctime, tzset: convert date/ ctime, localtime, gmtime, ctime.3c
 cu: call another UNIX System. cu.1c
 ttt: cubic: tic-tac-toe. ttt.6
 gethostname: get name of current host. gethostname.2
 hostname: set or print name of current host system hostname.1
 activity. sact: print current SCCS file editing sact.1
 uname: print name of current UNIX System. uname.1
 uname: get name of current UNIX system. uname.2
 slot in the utmp file of the current user. /find the ttyslot.3c
 getcwd: get pathname of the current working directory. getcwd.3c
 spline: interpolate smooth curve. spline.1g
 name of the user. cuserid: get character login cuserid.3s
 of each line of a file. cut: cut out selected fields cut.1
 each line of a file. cut: cut out selected fields of cut.1
 constant-width text for/ cw, checkcw: prepare cw.1
 cross reference. cxref: generate C program cxref.1
 cron: clock daemon. cron.1m
 errdemon: error-logging daemon. errdemon.1m
 terminate the error-logging daemon. errstop: errstop.1m
 runacct: run daily accounting. runacct.1m
 backup. filesave, tapesave: daily/weekly UNIX file system filesave.1m
 /handle special functions of DASI 300 and 300s terminals. 300.1
 special functions of the DASI 450 terminal. /handle 450.1
 blt, blt512: block transfer data. blt.3
 prof: display profile data. prof.1
 /time a command; report process data and system activity. timex.1
 termcap: terminal capability port. ttytype: termcap.5
 ttytype: data base of terminal types by ttytype.4
 /sget: access long numeric data in a machine independent/ sputl.3x
 plock: lock process, text, or data in memory. plock.2
 call. stat: data returned by stat system stat.5
 brk, sbrk: change data segment space allocation. brk.2
 types: primitive system data types. types.5

join: relational database operator. join.1
 udp: Internet User Datagram Protocol. udp.5
 date: print and set the date. date.1
 /asctime, tzset: convert date and time to string. ctime.3c
 date: print and set the date. date.1
 dc: desk calculator. dc.1
 optimal access time. dcopy: copy file systems for dcopy.1m
 dd: convert and copy a file. dd.1
 adb: debugger. adb.1
 fsdb: file system debugger. fsdb.1m
 eqnchar: special character definitions for eqn and neqn. eqnchar.5
 netmailer: deliver mail to. netmailer.8
 people. delivermail: deliver mail to arbitrary delivermail.8
 names. basename, dirname: deliver portions of path basename.1
 file. tail: deliver the last part of a tail.1
 aliases: aliases file for delivermail. aliases.7
 arbitrary people. delivermail: deliver mail to delivermail.8
 delta commentary of an SCCS delta. cdc: change the cdc.1
 file. delta: make a delta (change) to an SCCS delta.1
 delta. cdc: change the delta commentary of an SCCS cdc.1
 rmdel: remove a delta from an SCCS file. rmdel.1
 to an SCCS file. delta: make a delta (change) delta.1
 comb: combine SCCS deltas. comb.1
 msg: permit or deny messages. msg.1
 tbl, and eqn constructs. deroff: remove nroff/troff, deroff.1
 setkey, encrypt: generate DES encryption. crypt, crypt.3c
 close: close a file descriptor. close.2
 dup: duplicate an open file dup.2
 dc: desk calculator. dc.1
 file. access: determine accessibility of a access.2
 file: determine file type. file.1
 errors in the specified device. /on/off the extended exterr.1
 ioctl: control device. ioctl.2
 master: master device information table. master.4
 devnm: device name. devnm.1m
 devnm: device name. devnm.1m
 blocks. df: report number of free disk df.1m
 check and interactive/ fsck, fsck: file system consistency fsck.1m
 terminal line connection. dial: establish an out-going dial.3c
 bdiff: big diff. bdiff.1
 comparator. diff: differential file diff.1
 comparison. diff3: 3-way differential file diff3.1
 sdiff: side-by-side difference program. sdiff.1
 diffmk: mark differences between files. diffmk.1
 diff: differential file comparator. diff.1
 diff3: 3-way differential file comparison. diff3.1
 diffmk: mark differences between files. diffmk.1
 dir: format of directories. dir.4
 dircmp: directory comparison. dircmp.1
 directories. dir: format of dir.4
 directories. ls.1
 directories. rm.1
 directories. /count characters sumdir.1
 directory. cd.1
 directory. chdir.2
 directory. chroot.2
 directory. getcwd: get getcwd.3c
 directory. mkdir.1
 directory. mvdir.1m
 directory (Berkeley version). ls7.1
 directory clean-up. uclean.1m
 directory comparison. dircmp.1
 directory entry. unlink.2
 directory for a command. chroot.1m

/make a lost+found directory for fsck. mklost+fnd.1m
 pwd: working directory name. pwd.1
 ordinary file. mknod: make a directory, or a special or mknod.2
 path names. basename, dirname: deliver portions of basename.1
 printers. enable, disable: enable/disable LP enable.1
 acct: enable or disable process accounting. acct.2
 type, modes, speed, and line discipline. /set terminal getty.1m
 diskformat - format a disk. diskformat.1m
 sadp: disk access profiler. sadp.1
 df: report number of free disk blocks. df.1m
 disk tune - tune floppy disk settling time parameters. disk tune.1m
 du: summarize disk usage. du.1
 diskformat - format a disk. diskformat.1m
 disk tune - tune floppy disk settling time parameters. disk tune.1m
 mount, umount: mount and dismount file system. mount.1m
 rain: animated raindrops display. rain.6
 /view: screen oriented (visual) display editor based on ex. vi.1
 prof: display profile data. prof.1
 worms: animate worms on a display terminal. worms.6
 hypot: Euclidean distance function. hypot.3m
 /lcong48: generate uniformly distributed pseudo-random/ documents. mm: the MM mm.5
 macro package for formatting documents. /the OSDD adapter mosd.5
 macro package for formatting documents formatted with the/ mm.1
 mm, osdd, checkmm: print/check documents, view graphs, and mmt.1
 nulladm,/ chargefee, ckpacct, dodisk, lastlogin, monacct, acctsh.1m
 whodo: who is doing what. whodo.1m
 suitable for Motorola S-record downloading. /ASCII formats hex.1
 /Motorola S-records from downloading into a file. rcvhex.1
 nrand48, mrand48, jrand48, drand48, erand48, lrand48, drand48.3c
 arithmetic: provide drill in number facts. arithmetic.6
 du: summarize disk usage. du.1
 extract error records from dump. errdead: errdead.1m
 od: octal dump. od.1
 an object file. dump: dump selected parts of dump.1
 object file. dump: dump selected parts of an dump.1
 descriptor. dup: duplicate an open file dup.2
 descriptor. dup: duplicate an open file dup.2
 The alien invaders attack the earth. aliens: aliens.6
 echo: echo arguments. echo.1
 echo: echo arguments. echo.1
 floating-point number to/ ecvt, fcvt, gcvt: convert ecvt.3c
 program. end, etext, ed, red: text editor. ed.1
 ed, red: text editor. ed.1
 editor. ex.1
 editor. ld.1
 editor. sed.1
 oriented (visual) display editor based on ex. /screen vi.1
 se: screen editor for video terminals. se.1
 a.out: assembler and link editor output. a.out.4
 /user, real group, and effective group IDs. getuid.2
 and/ /getegid: get real user, effective user, real group, getuid.2
 Language. efl: Extended Fortran efl.1
 split fortran, ratfor, or efl files. fsplit: fsplit.1
 for a pattern. grep, enable/disable LP printers. grep.1
 enable/disable LP printers. enable, disable: enable.1
 accounting. acct: enable or disable process acct.2
 enable, disable: enable/disable LP printers. enable.1
 crypt: encode/decode. crypt.1
 encryption. crypt, setkey, encrypt: generate DES crypt.3c
 setkey, encrypt: generate DES encryption. crypt, crypt.3c

makekey: generate encryption key. makekey.1
 locations in program. end.3c
 /getgrgid, getgrnam, setgrent, endgrent: get group file/ getgrent.3c
 socket: create an endpoint for communication. socket.2
 /getpwuid, getpwnam, setpwent, endpwent: get password file/ getpwent.3c
 utmp/ /pututline, setutent, endutent, utmpname: access getut.3c
 convert Arabic numerals to English. number: number.6
 nlist: get entries from name list. nlist.3c
 man, manprog: print entries in this manual. man.1
 man: macros for formatting entries in this manual. man.5
 endgrent: get group file entry. /getgrnam, setgrent, getgrent.3c
 endpwent: get password file entry. /getpwnam, setpwent, getpwent.3c
 utmpname: access utmp file entry. /setutent, endutent, getut.3c
 putpwent: write password file entry. putpwent.3c
 unlink: remove directory entry. unlink.2
 utmp, wtmp: utmp and wtmp entry formats. utmp.4
 command execution.
 env: set environment for env.1
 environ: user environment. environ.4
 environ: user environment. environ.5
 environ: user environment. environ.4
 environ: user environment. environ.5
 environment. environ.5
 environment. printenv.1
 environment at login time. profile.4
 environment for command env.1
 environment name. getenv.3c
 eqn and neqn. /special eqnchar.5
 eqn constructs. deroff: deroff.1
 eqn, neqn, checkeq: format eqn.1
 eqnchar: special character eqnchar.5
 erand48, jrand48, / drand48, drand48.3c
 complementary error function. erf.3m
 complementary error/ erf, erf.3m
 err: error-logging interface. err.7
 errdead: extract error records errdead.1m
 daemon. errdemon.1m
 format. errfile.4
 system error/ perror, perror.3c
 function and complementary erf.3m
 complementary/ erf, erfc: erf.3m
 error function and erf.3m
 error function and erf.3m
 error message file by mkstr.1
 massaging C/ mkstr: create an perror.3c
 sys_errlist, sys_nerr: system intro.2
 to system calls and
 errdead: extract error records errdead.1m
 from dump.
 matherr: error-handling function. matherr.3m
 errfile: error-log file format. errfile.4
 errdemon: error-logging daemon. errdemon.1m
 errstop: terminate the errstop.1m
 err: error-logging interface. err.7
 process a report of logged errprt.1m
 hashcheck: find spelling spell.1
 /- turn on/off the extended exterr.1
 logged errors. errprt.1m
 error-logging daemon. errstop.1m
 robots. autorobots: autorobots.6
 robots: Escape from the automatic robots.6
 chase: Try to escape the killer robots. chase.6
 terminal line/ dial: dial.3c
 setmnt: establish mount table. setmnt.1m
 bnet.
 in program. end, hosts.7
 etext, edata: last locations end.3c
 hypot: Euclidean distance function. hypot.3m
 expression. expr: evaluate arguments as an expr.1
 test: condition test.1
 display editor based on ex. /screen oriented (visual) vi.1

ex, edit: text editor ex.1
 exclusive file regions for lockf.2
 execl, execl, execl, execl, execl, exec.2
 execl, execl, execl, exec.2
 execl, execl, execl, exec.2
 execute a file. /execle, exec.2
 execute command. xargs: xargs.1
 execute commands at a later at.1
 execute regular expression. regcmp.3x
 execution. env: env.1
 execution. uux.1c
 execution for an interval. sleep.1
 execution for an interval. sleep.3c
 execution profile. monitor.3c
 execution time profile. profil.2
 execv, execl, execl, execl, exec.2
 execl, execl, execl, exec.2
 /execv, execl, execl, execl, exec.2
 execl, execl, execl, exec.2
 system calls. link, unlink: link.1m
 a new file or rewrite an creat.2
 process. exit.2
 exit, exit.2
 _exit: terminate process. exit.2
 exponential, logarithm,/ exp.3m
 pcat, unpack: compress and pack.1
 adventure: an adventure.6
 exp, log, log10, pow, sqrt: exp.3m
 exponential, logarithm, power,/ exp.3m
 expression. expr.1
 expr: evaluate arguments as an expr.1
 expression. regcmp, regex: regcmp.3x
 expression compile. regcmp.1
 expression compile and match regexp.5
 extended errors in the/ exterr.1
 Extended Fortran Language. efl.1
 extended TTY-37 type-box. greek.5
 exterr - turn on/off the exterr.1
 extract error records from errdead.1m
 remainder,/ floor, ceil, fmod, floor.3m
 fabs: floor, ceiling, floor.3m
 factor: factor.1
 factor: factor a number. factor.1
 true, true.1
 false: provide truth values. true.1
 data in a machine independent sputl.3x
 fnc: /access long numeric fnc.1m
 fast incremental backup. abort.3c
 abort: generate an IOT abort.3c
 a stream. fclose.3s
 floating-point number/ fcvt.3c
 ecvt, gcvt: convert fcvt.3c
 fdopen: open a stream. fdopen.3s
 fopen, freopen, fopen.3s
 status inquiries. ferror, ferror.3s
 feof, clearerr, fileno: stream ferror.3s
 ferror, feof, clearerr, ferror.3s
 ff: list file names and ff.1m
 stream. fclose, fclose.3s
 word from/ getc, getchar, getc.3s
 stream. gets, gets.3s
 stream. gets, gets.3s
 pattern. grep, egrep, grep.1
 fgrep: search a file for a grep.1
 determine accessibility of a access.2
 chmod: change mode of access.2
 change owner and group of a chown.2
 core: format of core image core.4
 fields of each line of a cut.1
 dd: convert and copy a dd.1
 a delta (change) to an SCCS delta.1
 selected parts of an object dump.1
 execl, execl, execl, exec.2
 file. access: access.2
 file. chown.2
 file. chown: chown.2
 file. core.4
 file. cut: cut out selected cut.1
 file. dd.1
 file. delta: make delta.1
 file. dump: dump dump.1
 file. /execv, execl, execl, exec.2

on character frequencies in a	file. freq: report	freq.1	a file system. ff: list	file names and statistics for	ff.1m
get: get a version of an SCCS	file.	get.1	/find the slot in the utmp	file of the current user.	ttyslot.3c
group: group	file.	group.4	put: puts a	file onto a remote machine.. . . .	put.1c
issue: issue identification	file.	issue.4	put7: puts a	file onto a remote machine.. . . .	put7.1c
link: link to a	file.	link.2	/identify processes using a	file or file structure.	fuser.1m
mknod: build special	file.	mknod.1m	one. creat: create a new	file or rewrite an existing	creat.2
or a special or ordinary	file. /make a directory,	mknod.2	viewing. more:	file perusal filter for crt	more.1
change the format of a text	file. newform:	newform.1	lseek: move read/write	file pointer.	lseek.2
null: the null	file.	null.7	/rewind, ftell: reposition a	file pointer in a stream.	fseek.3s
passwd: password	file.	passwd.4	lockf: provide exclusive	file regions for reading or/	lockf.2
or subsequent lines of one	file. /lines of several files	paste.1	bfs: big	file scanner.	bfs.1
prs: print an SCCS	file.	prs.1	stat, fstat: get	file status.	stat.2
from downloading into a	file. /Motorola S-records	rcvhex.1	processes using a file or	file structure. /identify	fuser.1m
read: read from	file.	read.2	names and statistics for a	file system. ff: list file	ff.1m
remove a delta from an SCCS	file. rmdel:	rmdel.1	mkfs: construct a	file system.	mkfs.1m
two versions of an SCCS	file. sccsdiff: compare	sccsdiff.1	mkfs512: construct a	file system.	mkfs512.1m
sccsfile: format of SCCS	file.	sccsfile.4	umount: mount and dismount	file system. mount,	mount.1m
size: size of an object	file.	size.1	mount: mount a	file system.	mount.2
in an object, or other binary	file. /the printable strings	strings.1	umount: unmount a	file system.	umount.2
checksum and block count of a	file. sum: print	sum.1	tapesave: daily/weekly UNIX	file system backup. filesave,	filesave.1m
sum and count blocks in a	file. sum7:	sum7.1	and interactive/ fsck, dfsc:	file system consistency check	fsck.1m
deliver the last part of a	file. tail:	tail.1	fsdb:	file system debugger.	fsdb.1m
tmpfile: create a temporary	file.	tmpfile.3s	volume.	file system: format of system	fs.4
create a name for a temporary	file. tmpnam, tmpnam:	tmpnam.3s	ustat: get	file system statistics.	ustat.2
and modification times of a	file. touch: update access	touch.1	mnttab: mounted	file system table.	mnttab.4
undo a previous get of an SCCS	file. unget:	unget.1	access time. dcopy: copy	file systems for optimal	dcopy.1m
report repeated lines in a	file. uniq:	uniq.1	fsck. checklist: list of	file systems processed by	checklist.4
val: validate SCCS	file.	val.1	volcopy, labelit: copy	file systems with label/	volcopy.1m
write: write on a	file.	write.2	ftw: walk a	file tree.	ftw.3c
times. utime: set	file access and modification	utime.2	file: determine	file type.	file.1
hpio: HP 2645A terminal tape	file archiver.	hpio.1	umask: set	file-creation mode mask.	umask.1
tar: tape	file archiver.	tar.1	error, feof, clearerr,	fileno: stream status/	ferror.3s
cpio: copy	file archives in and out.	cpio.1	and print process accounting	file(s). acctcom: search	acctcom.1
mkstr: create an error message	file by massaging C source.	mkstr.1	merge or add total accounting	files. acctmerg:	acctmerg.1m
pwck, grpck: password/group	file checkers.	pwck.1m	create and administer SCCS	files. admin:	admin.1
diff: differential	file comparator.	diff.1	cat: concatenate and print	files.	cat.1
diff3: 3-way differential	file comparison.	diff3.1	cmp: compare two	files.	cmp.1
fcntl:	file control.	fcntl.2	lines common to two sorted	files. comm: select or reject	comm.1
fcntl:	file control options.	fcntl.5	cp, ln, mv: copy, link or move	files.	cp.1
rcp: remote	file copy	rcp.1	mark differences between	files. diffmk:	diffmk.1
UNIX System-to-UNIX System	file copy. /uupick: public	uuto.1c	find: find	files.	find.1
umask: set and get	file creation mask.	umask.2	format specification in text	files. fspec:	fspec.4
close: close a	file descriptor.	close.2	fortran, ratfor, or efl	files. fsplit: split	fsplit.1
dup: duplicate an open	file descriptor.	dup.2	string, format of graphical	files. /graphical primitive	gps.4
sact: print current SCCS	file: determine file type.	file.1	intro: introduction to special	files.	intro.7
setgrent, endgrent: get group	file editing activity.	sact.1	unpack: compress and expand	files. pack, pcat,	pack.1
endpwent: get password	file entry. /getgrnam,	getgrent.3c	pr: print	files.	pr.1
utmpname: access utmp	file entry. /setpwent,	getpwent.3c	sort: sort and/or merge	files.	sort.1
putpwent: write password	file entry. /endutent,	getut.3c	reports version number of	files. version:	version.1
ctags: maintain a tags	file entry.	putpwent.3c	what: identify SCCS	files.	what.1
grep, egrep, fgrep: search a	file for a C program.	ctags.1	updater: update	files between two machines.	updater.1
aliases: aliases	file for a pattern.	grep.1	updater: update	files between two machines.	updater.1m
acct: per-process accounting	file for delivermail.	aliases.7	freq: recover	files from a backup tape.	freq.1m
ar: archive (library)	file format.	acct.4	and count characters in the	files in the given/ /sum	sumdir.1
errfile: error-log	file format.	ar.4	hex: translates object	files into ASCII formats/	hex.1
pnch:	file format.	errfile.4	rm, rmdir: remove	files or directories.	rm.1
intro: introduction to	file format for card images.	pnch.4	/merge same lines of several	files or subsequent lines of/	paste.1
take: takes a	file formats.	intro.4	daily/weekly UNIX file system/	filesave, tapesave:	filesave.1m
take7: takes a	file from a remote machine.	take.1c	greek: select terminal	filter.	greek.1
see: see what a	file from a remote machine.. . . .	take7.1c	nl: line numbering	filter.	nl.1
split: split a	file has in it.	see.1	more: file perusal	filter for crt viewing.	more.1
mktemp: make a unique	file into pieces.	split.1	col:	filter reverse line-feeds.	col.1
ctermid: generate	file name.	mktemp.3c	tplot: graphics	filters.	tplot.1g
	file name for terminal.	ctermid.3s		finc: fast incremental backup.	finc.1m

find: find files. find.1
 find: find files. find.1
 hyphen: find hyphenated words. hyphen.1
 ttyname, isatty: find name of a terminal. ttyname.3c
 object library. lorder: find ordering relation for an
 hashmake, spellin, hashcheck: find spelling errors. spell, spell.1
 an object, or other/ strings: find the printable strings in strings.1
 of the current user. ttyslot: find the slot in the utmp file ttyslot.3c
 fish: play "Go Fish". fish.6
 fish: play "Go Fish". fish.6
 (sh only). nohup: run nohup.1
 fitting. tee.1
 floating-point number. atof.3c
 floating-point number to/ ecvt.3c
 floating-point numbers. frexp.3c
 floor, ceil, fmod, fabs: floor, ceil, fmod, fabs: floor.3m
 floor, ceiling, remainder,/ floor.3m
 floppy disk settling time disktime.1m
 flow graph. cflow.1
 flush a stream. fclose.3s
 fmod, fabs: floor, ceiling, floor.3m
 fopen, freopen, fdopen: open a fopen.3s
 fork: create a new process. fork.2
 format. acct: acct.4
 format. ar.4
 format. errfile.4
 format. tp.4
 format a disk. diskformat.1m
 format for card images. pnch.4
 format mathematical text for eqn.1
 format of a text file. newform.1
 format of an inode. inode.4
 core: format of core image file. core.4
 cpio: format of cpio archive. cpio.4
 dir: format of directories. dir.4
 /graphical primitive string, gps.4
 sccsfile: format of SCCS file. sccsfile.4
 file system: format of system volume. fs.4
 files. fspec: format specification in text fspec.4
 troff. tbl: format tables for nroff or tbl.1
 nroff: format text. nroff.1
 formats. intro.4
 formats. utmp, utmp.4
 formats suitable for Motorola/ hex.1
 formatted input. scanf.3s
 formatted output. printf, printf.3s
 formatted with the MM macros. mm.1
 formatting a permuted index. mptx.5
 formatting and typesetting. nroff7.1
 formatting and typesetting. troff7.1
 formatting documents. mm.5
 formatting documents. /the mosd.5
 formatting entries in this man.5
 FORTRAN compiler. fortran.1
 fortran: FORTRAN compiler. fortran.1
 Fortran Language. efl.1
 fortran, ratfor, or efl fsplit.1
 fortune: print a random, fortune.6
 fprintf, sprintf: print printf.3s
 fputc, putw: put character or putc.3s
 fputs: put a string on a puts.3s
 fread, fwrite: binary fread.3s
 frec: recover files from a frec.1m
 free disk blocks. df.1m

memory allocator. malloc, free, realloc, calloc: main malloc.3c
 stream. fopen, freopen, fdopen: open a fopen.3s
 frequencies in a file. freq: report on character freq.1
 freq: report on character frequencies in a file. freq.1
 parts of floating-point/ frexp, ldexp, modf: manipulate frexp.3c
 frec: recover files from a backup tape. frec.1m
 take: takes a file from a remote machine. take.1c
 take7: takes a file from a remote machine.. take7.1c
 receive: receive message from a socket. receive.2
 send: send message from a socket. send.2
 gets, fgets: get a string from a stream. gets.3s
 rmdel: remove a delta from an SCCS file. rmdel.1
 getopt: get option letter from argument vector. getopt.3c
 /translates Motorola S-records from downloading into a file. rcvhex.1
 errdead: extract error records from dump. errdead.1m
 read: read from file. read.2
 ncheck: generate names from i-numbers. ncheck.1m
 nlist: get entries from name list. nlist.3c
 acctcms: command summary from per-process accounting/ acctcms.1m
 getw: get character or word from stream. /getchar, fgets, getc.3s
 autorobots: Escape from the automatic robots. autorobots.6
 robots: Escape from the robots. robots.6
 getpw: get name from UID. getpw.3c
 formatted input. scanf, fsck, sscanf: convert scanf.3s
 of file systems processed by fsck. checklist: list checklist.4
 a lost + found directory for fsck. mklost + found: make mklost + fnd.1m
 consistency check and/ fsck, dfck: file system fsck.1m
 fsdb: file system debugger. fsdb.1m
 reposition a file pointer in/ fseek, rewind, ftell: fseek.3s
 text files. fspec: format specification in fspec.4
 or efl files. fsplit: split fortran, ratfor, fsplit.1
 stat. fstat: get file status. stat.2
 pointer in a/ fseek, rewind, ftell: reposition a file fseek.3s
 ftw: walk a file tree. ftw.3c
 and complementary error function. /error function erf.3m
 gamma: log gamma function. gamma.3m
 hypot: Euclidean distance function. hypot.3m
 matherr: error-handling function. matherr.3m
 error/ erf, erf: error function and complementary erf.3m
 j0, j1, jn, y0, y1, yn: Bessel functions. bessell.3m
 logarithm, power, square root functions. /sqrt: exponential, exp.3m
 remainder, absolute value functions. /floor, ceiling, floor.3m
 sinh, cosh, tanh: hyperbolic functions. sinh.3m
 atan, atan2: trigonometric functions. /tan, asin, acos, trig.3m
 300, 300s: handle special functions of DASI 300 and 300s/ 300.1
 hp: handle special functions of HP 2640 and/ hp.1
 terminal. 450: handle special functions of the DASI 450 450.1
 using a file or file/ fuser: identify processes fuser.1m
 fread, fwrite: binary input/output. fread.3s
 connect accounting records. fwtmp, wtmpfix: manipulate fwtmp.1m
 adventure: an exploration game. adventure.6
 moo: guessing game. moo.6
 trek: trekkie game. trek.6
 worm: Play the growing worm game. worm.6
 cribbage: the card game cribbage. cribbage.6
 back: the game of backgammon. back.6
 bj: the game of black jack. bj.6
 craps: the game of craps. craps.6
 wump: the game of hunt-the-wumpus. wump.6
 life: play the game of life. life.6
 intro: introduction to games. intro.6
 gamma: log gamma function. gamma.3m
 gamma: log gamma function. gamma.3m
 number to string. ecvt, fcvt, gcvt: convert floating-point ecvt.3c

maze: generate a maze. maze.6
 abort: generate an IOT fault. abort.3c
 cflow: generate C flow graph. cflow.1
 reference. cxref: generate C program cross cxref.1
 crypt, setkey, encrypt: generate DES encryption. crypt.3c
 makekey: generate encryption key. makekey.1
 terminal. ctermid: generate file name for ctermid.3s
 ncheck: generate names from i-numbers. ncheck.1m
 lexical tasks. lex: generate programs for simple lex.1
 /srand48, seed48, lcong48: generate uniformly distributed/ drand48.3c
 srand: simple random-number generator. rand, rand.3c
 gets, fgets: get a string from a stream. gets.3s
 get: get a version of an SCCS file. get.1
 ulimit: get and set user limits. ulimit.2
 the user. cuserid: get character login name of cuserid.3s
 getc, getchar, fgetc, getw: get character or word from/ getc.3s
 nlist: get entries from name list. nlist.3c
 umask: set and get file creation mask. umask.2
 stat, fstat: get file status. stat.2
 ustat: get file system statistics. ustat.2
 file. get: get a version of an SCCS get.1
 /getgrnam, setgrent, endgrent: get group file entry. getgrent.3c
 getlogin: get login name. getlogin.3c
 logname: get login name. logname.1
 msgget: get message queue. msgget.2
 getpw: get name from UID. getpw.3c
 gethostname: get name of current host. gethostname.2
 system. uname: get name of current UNIX uname.2
 unget: undo a previous get of an SCCS file. unget.1
 argument vector. getopt: get option letter from getopt.3c
 /getpwnam, setpwnent, endpwnent: get password file entry. getpwnent.3c
 working directory. getcwd: get pathname of current getcwd.3c
 times. times: get process and child process times.2
 and/ getpid, getpgrp, getppid: get process, process group, getpid.2
 /geteuid, getgid, getegid: get real user, effective user,/ geteuid.2
 semget: get set of semaphores. semget.2
 shmget: get shared memory segment. shmget.2
 tty: get the terminal's name. tty.1
 time: get time. time.2
 get character or word from/ getc, getchar, fgetc, getw: getc.3s
 character or word from/ getc, getchar, fgetc, getw: get getc.3s
 current working directory. getcwd: get pathname of getcwd.3c
 getuid, geteuid, getgid, getegid: get real user,/ getuid.2
 environment name. getenv: return value for getenv.3c
 real user, effective/ getuid, geteuid, getgid, getegid: get getuid.2
 user,/ getuid, geteuid, getgid, getegid: get real getuid.2
 setgrent, endgrent: get group/ getgrent, getgrgid, getgrnam, getgrent.3c
 endgrent: get group/ getgrent, getgrgid, getgrnam, setgrent, endgrent: getgrent.3c
 get group/ getgrent, getgrgid, getgrnam, setgrent, endgrent: getgrent.3c
 current host. gethostname: get name of gethostname.2
 getlogin: get login name. getlogin.3c
 argument vector. getopt: get option letter from getopt.3c
 getopt: parse command options. getopt.1
 getpass: read a password. getpass.3c
 getpgrp, getppid: get process, process group, and/ getpid.2
 process group, and/ getpid, getppid: get process, process group, and/ getpid.2
 group, and/ getpid, getpgrp, getpw: get name from UID. getpw.3c
 setpwnent, endpwnent: get/ getpwnam, setpwnent, endpwnent: getpwnent.3c
 get/ getpwnent, endpwnent: get/ getpwnent, getpwnent.3c
 a stream. gets, fgets: get a string from gets.3s
 and terminal settings used by getty. gettydefs: speed gettydefs.4
 modes, speed, and line/ getty: set terminal type, getty.1m

ct: spawn getty to a remote terminal. ct.1c
 settings used by getty. gettydefs: speed and terminal gettydefs.4
 getegid: get real user,/ pututline, setutent,/ getuid.2
 pututline, setutent,/ getut.3c
 setutent, endutent,/ getut.3c
 getutent, getutid, getutline, getutid, getutline, pututline, getut.3c
 setutent,/ getut.3c
 from/ getc, getchar, fgetc, getutline, pututline, getut.3c
 convert/ ctime, localtime, getw: get character or word getc.3s
 fish: play gmtime, asctime, tzset: ctime.3c
 "Go Fish". fish.6
 setjmp, longjmp: non-local goto. setjmp.3c
 string, format of graphical/ gps: graphical primitive gps.4
 cflow: generate C flow graph. cflow.1
 sag: system activity graph. sag.1g
 primitive string, format of graphical files. /graphical gps.4
 format of graphical/ gps: graphical primitive string, gps.4
 tplot: graphics filters. tplot.1g
 TTY-37 type-box. greek: graphics for the extended greek.5
 plot: graphics interface. plot.4
 subroutines. plot: graphics interface plot.3x
 mv: typeset documents, view graphs, and slides. mmt, mmt.1
 package for typesetting view graphs and slides. /macro mv.5
 extended TTY-37 type-box. greek: graphics for the greek.5
 file for a pattern. greek: select terminal filter. greek.1
 chown, chgrp: change owner or grep, egrep, fgrep: search a grep.1
 newgrp: log in to a new group. chown.1
 /user, effective user, real group. newgrp.1
 /getppid: get process, process group, and effective group/ getuid.2
 group, and parent process IDs. getppid.2
 group file. group.4
 setgrent, endgrent: get group file entry. /getgrnam, getgrent.3c
 group: group file. group.4
 setpgrp: set process group ID. setpgrp.2
 real group, and effective group IDs. /effective user, getuid.2
 setuid, setgid: set user and group IDs. setuid.2
 id: print user and group IDs and names. id.1
 chown: change owner and group of a file. chown.2
 a signal to a process or a group of processes. /send kill.2
 update, and regenerate groups of programs. /maintain, make.1
 worm: Play the growing worm game. worm.6
 checkers. pwck, grpck: password/group file pwck.1m
 ssignal: software signals. ssignal.3c
 hangman: guess the word. hangman.6
 moo: guessing game. moo.6
 DASI 300 and 300s/ 300, 300s: handle special functions of 300.1
 2640 and 2621-series/ hp: handle special functions of HP hp.1
 the DASI 450 terminal. 450: handle special functions of 450.1
 information for bad block handling. /alternate block altblk.4
 hangman: guess the word. hangman.6
 hangups (sh only). nohup.1
 nohup: run a command immune to hcreate, hdestroy: manage hsearch.3c
 hcreate, hdestroy: manage hash spell, hashmake, spellin, spell.1
 find spelling errors. spell, hashcheck: find spelling/ spell.1
 search tables. hsearch, hashmake, spellin, hashcheck: hsearch.3c
 tables. hsearch, hcreate, hdestroy: manage hash search hsearch.3c
 help: ask for help. help.1
 help: ask for help. help.1
 into ASCII formats suitable/ hex: translates object files hex.1
 fortune: print a random, hopefully interesting, adage. fortune.6
 get name of current host. gethostname: gethostname.2
 sethostname: set name of host cpu. sethostname.2
 ruptime: show host status of local machines ruptime.1
 set or print name of current host system hostname: hostname.1
 /etc/hosts: host table for bnet. hosts.7
 current host system hostname: set or print name of hostname.1

rhost, raddr: look up internet
 handle special functions of archiver. hpio: HP 2645A terminal tape file archiver. hpio: HP 2645A terminal tape manage hash search tables. wump: the game of sinh, cosh, tanh:
 hyphen: find function.
 semaphore set or shared memory setpgrp: set process group and names. issue: issue file or file/ fuser: what: group, and parent process group, and effective group setgid: set user and group id: print user and group core: format of core pnch: file format for card only). nohup: run a command finc: fast
 long numeric data in a machine /tgoto, tputs: terminal for formatting a permuted ptx: permuted and teletypes. last: family. inittab: script for the initialization. init, telinit: process control /rc, powerfail: system socket. connect: process. popen, pclose: process. clri: clear inode: format of an sscanf: convert formatted push character back into fread, fwrite: binary stdout: standard buffered fileno: stream status uustat: uucp status install:
 atol, atoi: convert string to abs: return /l64a: convert between long 3-byte integers and long /ltol3: convert between 3-byte bcopy: system consistency check and print a random, hopefully err: error-logging loop: software loopback plot: graphics termio: general terminal tty: controlling terminal rhost, raddr: look up hosts by name or address. rhost.3 HP 2640 and 2621-series/ hp: hp.1 HP 2645A terminal tape file hpio.1 hp: handle special functions hp.1 hpio: HP 2645A terminal tape hpio.1 hsearch, hcreate, hdestroy: hsearch.3c hunt-the-wumpus. wump.6 hyperbolic functions. sinh.3m hyphen: find hyphenated words. hyphen.1 hyphenated words. hyphen.1 hypot: Euclidean distance hypot.3m id. /remove a message queue, ipcrm.1 ID. setpgrp.2 id: print user and group IDs id.1 identification file. issue.4 identify processes using a fuser.1m identify SCCS files. what.1 IDs. /get process, process getpid.2 IDs. /effective user, real getuid.2 IDs. setuid, setuid.2 IDs and names. id.1 image file. core.4 images. pnch.4 immune to hangups (sh nohup.1 incremental backup. finc.1m independent fashion. /access sputl.3x independent operation/ termcap.3 index. /the macro package mptx.5 index. ptx.1 indicate last logins of users last.1 inet: Internet protocol inet.5 init process. inittab.4 init, telinit: process control init.1m initialization. init.1m initialization shell scripts. brc.1m initiate a connection on a connect.2 initiate pipe to/from a popen.3s inittab: script for the init inittab.4 i-node. clri.1m inode. inode.4 inode: format of an inode. inode.4 input. scanf, fscanf, scanf.3s input stream. ungetc: ungetc.3s input/output. fread.3s input/output package. stdout.3s inquiries. /feof, clearerr, ferror.3s inquiry and job control. uustat.1c install commands. install.1m install: install commands. install.1m integer. strtol, strtol.3c integer absolute value. abs.3c integer and base-64 ASCII/ a64l.3c integers. /convert between l3tol.3c integers and long integers. l3tol.3c interactive block copy. bcopy.1m interactive repair. /file fsck.1m interesting. adage. fortune: fortune.6 interface. err.7 interface. lo.5 interface. plot.4 interface. termio.7 interface. tty.7 interface subroutines. plot.3x internet hosts by name or/ rhost.3

ip: Internet Protocol. ip.5 inet: Internet protocol family. inet.5 Protocol. tcp: Internet Transmission Control tcp.5 Protocol. udp: Internet User Datagram udp.5 spline: interpolate smooth curve. spline.1g characters. asa: interpret ASA carriage control asa.1 sno: SNOBOL interpreter. sno.1 interpreter) with C-like csh.1 interprocess channel. pipe.2 inter-process communication ipc.1 interprocess communication stdipc.3c interval. sleep: sleep.1 interval. sleep.3c intro: introduction to intro.1 intro: introduction to file intro.4 intro: introduction to games. intro.6 intro: introduction to intro.5 intro: introduction to special intro.7 intro: introduction to intro.3 intro: introduction to system intro.2 intro: introduction to system intro.1m intro: introduction to system intro.8 intro: introduction to commands and intro.1 intro: introduction to file formats. intro.4 intro: introduction to games. intro.6 intro: introduction to miscellany. intro.5 facilities. net: introduction to networking net.5 intro: introduction to special files. intro.7 and libraries. intro: introduction to subroutines intro.3 and error numbers. intro: introduction to system calls intro.2 maintenance commands/ intro: introduction to system intro.1m maintenance/ intro: introduction to system intro.8 ncheck: generate names from ncheck.1m aliens: The alien invaders attack the earth. aliens.6 select: synchronous i/o multiplexing. select.2 ioctl: control device. ioctl.2 IOT fault. abort.3c ip: Internet Protocol. ip.5 ipcrm: remove a message queue, ipcrm.1 ipc: report inter-process ipc.1 isalnum, isspace, ispunct./ ctype.3c isalpha, isupper, islower, ctype.3c isascii: classify characters. ctype.3c terminal. ttyname, isatty: find name of a ttyname.3c isctrl, isascii: classify/ ctype.3c isdigit, isxdigit, isalnum./ ctype.3c isgraph, isctrl, isascii:/ ctype.3c islower, isdigit, isxdigit, ctype.3c isprint, isgraph, isctrl,/ ctype.3c ispunct, isprint, isgraph,/ ctype.3c isspace, ispunct, isprint,/ ctype.3c system: issue a shell command. system.3s issue: issue identification file. issue.4 issue: issue identification issue.4 isupper, islower, isdigit, ctype.3c isxdigit, isalnum, isspace,/ ctype.3c news: print news news.1 functions. j0, j1, jn, y0, y1, yn: Bessel bessell.3m j1, jn, y0, y1, yn: Bessel bessell.3m jack. bj.6 jn, y0, y1, yn: Bessel bessell.3m operator. join: relational database join.1 /lrand48, nrand48, mrand48, jrand48, srand48, seed48,/ drand48.3c makekey: generate encryption key. makekey.1

killall: kill all active processes. killall.1m
 process or a group of/ kill: send a signal to a kill.2
 kill: terminate a process. kill.1
 processes. killall: kill all active killall.1m
 chase: Try to escape the killer robots. chase.6
 mem. kmem: core memory. mem.7
 quiz: test your knowledge. quiz.6
 3-byte integers and long/ l3tol, ltol3: convert between l3tol.3c
 integer and base-64/ a64l, l64a: convert between long a64l.3c
 copy file systems with label checking. /labelit: volcopy.1m
 with label checking. volcopy, labelit: copy file systems volcopy.1m
 scanning and processing language. awk: pattern awk.1
 arbitrary-precision arithmetic language. bc: bc.1
 efl: Extended Fortran Language. efl.1
 command programming language. /standard/restricted sh.1
 cpp: the C language preprocessor. cpp.1
 chargefee, ckpacct, dodisk, lastlogin, monacct, nulladm,/ /jrand48, srand48, seed48, lcong48: generate uniformly/ drand48.3c
 ld: link editor. ld.1
 of floating-point/ frexp, ldexp, modf: manipulate parts frexp.3c
 getopt: get option letter from argument vector. getopt.3c
 simple lexical tasks. lex: generate programs for lex.1
 generate programs for simple lexical tasks. lex: lex.1
 to subroutines and libraries. /introduction intro.3
 relation for an object library. /find ordering lorder.1
 ar: archive (library) file format. ar.4
 ar: archive and library maintainer. ar.1
 ulimit: get and set user limits. ulimit.2
 line: read one line. line.1
 an out-going terminal line connection. /establish dial.3c
 type, modes, speed, and line discipline. /set terminal getty.1m
 nl: line numbering filter. nl.1
 out selected fields of each line of a file. cut: cut cut.1
 send/cancel requests to an LP line printer. lp, cancel: lp.1
 lpr: line printer spooler. lpr.1
 line: read one line. line.1
 lsearch: linear search and update. lsearch.3c
 col: filter reverse line-feeds. col.1
 head: give first few lines. head.1
 files. comm: select or reject lines common to two sorted comm.1
 uniq: report repeated lines in a file. uniq.1
 of several files or subsequent lines of one file. /same lines paste.1
 subsequent/ paste: merge same lines of several files or paste.1
 link, unlink: exercise link and unlink system calls. link.1m
 ld: link editor. ld.1
 a.out: assembler and link editor output. a.out.4
 cp, ln, mv: copy, link: link to a file. link.2
 link or move files. cp.1
 link to a file. link.2
 and unlink system calls. link, unlink: exercise link link.1m
 lint: a C program checker. lint.1
 nlist: get entries from name list. nlist.3c
 nm: print name list. nm.1
 ls: list contents of directories. ls.1
 (Berkeley version). ls7: list contents of directory ls7.1
 for a file system. ff: list file names and statistics ff.1m
 by fsck. checklist: list of file systems processed checklist.4
 xargs: construct argument list(s) and execute command. xargs.1
 files. cp, ln, mv: copy, link or move cp.1
 tzset: convert date/ ctime, localtime, gmtime, asctime, ctime.3c
 end, etext, edata: last locations in program. end.3c
 memory. plock: lock process, text, or data in plock.2
 regions for reading or lockf: provide exclusive file lockf.2
 gamma: log gamma function. gamma.3m

newgrp: log in to a new group. newgrp.1
 exponential, logarithm,/ exp, log, log10, pow, sqrt: exp.3m
 logarithm, power,/ exp, log, log10, pow, sqrt: exponential, exp.3m
 /log10, pow, sqrt: exponential, logarithm, power, square root/ exp.3m
 errpt: process a report of logged errors. errpt.1m
 rwho: who is logged in on local machines rwho.1
 rlogin: remote login rlogin.1
 getlogin: get login name. getlogin.3c
 logname: get login name. logname.1
 cuserid: get character login name of the user. cuserid.3s
 logname: return login name of user. logname.3x
 passwd: change login password. passwd.1
 login: sign on. login.1
 setting up an environment at login time. profile: profile.4
 last: indicate last logins of users and teletypes. last.1
 logname: get login name. logname.1
 user. logname: return login name of logname.3x
 a64l, l64a: convert between long integer and base-64 ASCII/ a64l.3c
 between 3-byte integers and long integers. /ltol3: convert l3tol.3c
 sputl, sgetl: access long numeric data in a machine/ sputl.3x
 setjmp, longjmp: non-local goto. setjmp.3c
 interface. loop: software loopback lo.5
 loop: software loopback interface. lo.5
 for an object library. lorder: find ordering relation lorder.1
 mklost+fnd: make a lost+fnd directory for fsck. mklost+fnd.1m
 nice: run a command at low priority. nice.1
 requests to an LP line/ lp, cancel: send/cancel lp.1
 send/cancel requests to an LP line printer. lp, cancel: lp.1
 disable: enable/disable LP printers. enable, enable.1
 /lpshut, lpmove: start/stop the LP request scheduler and move/ lpsched.1m
 accept, reject: allow/prevent LP requests. accept.1m
 lpadmin: configure the LP spooling system. lpadmin.1m
 lpstat: print LP status information. lpstat.1
 request/ lpsched, lpshut, lpadmin: configure the LP lpadmin.1m
 spooling system. lpmove: start/stop the LP lpsched.1m
 lpr: line printer spooler. lpr.1
 start/stop the LP request/ lpsched, lpshut, lpmove: lpsched.1m
 LP request scheduler/ lpsched, lpshut, lpmove: start/stop the lpsched.1m
 information. lpstat: print LP status lpstat.1
 jrand48,/ drand48, erand48, lrand48, nrand48, mrand48, drand48.3c
 directories. ls: list contents of ls.1
 directory (Berkeley version). ls7: list contents of ls7.1
 update. lsearch: linear search and lsearch.3c
 pointer. lseek: move read/write file lseek.2
 integers and long/ l3tol, ltol3: convert between 3-byte l3tol.3c
 m4: macro processor. m4.1
 truth value about your/ m68k, pdp11, u3b, vax: provide machid.1
 put: puts a file onto a remote machine.. put.1c
 puts a file onto a remote machine.. put7: put7.1c
 takes a file from a remote machine. take: take.1c
 takes a file from a remote machine.. take7: take7.1c
 /access long numeric data in machine independent fashion.. sputl.3x
 show host status of local machines ruptime: ruptime.1
 who is logged in on local machines rwho: rwho.1
 update files between two machines. updater: updater.1
 update files between two machines. updater: updater.1m
 permuted index. mptx: the macro package for formatting a mptx.5
 documents. mm: the MM macro package for formatting mm.5
 mosd: the OSDD adapter macro package for formatting/ mosd.5
 view graphs and/ mv: a troff macro package for typesetting mv.5
 m4: macro processor. m4.1
 formatted with the MM macros. /print/check documents mm.1
 in this manual. man: macros for formatting entries man.5
 tp: magnetic tape format. tp.4

send mail to users or read
users or read mail.
netmail: the bnet network
netmailer: deliver
delivermail: deliver
mail, rmail: send
malloc, free, realloc, calloc:
program. ctags:
regenerate groups of/
ar: archive and library
intro: introduction to system
intro: introduction to system
SCCS file. delta:
mkdir:
or ordinary file. mknod:
for fsck. mklost+found:
mktemp:
regenerate groups of/
ssp:
banner:
key.
main memory allocator.
entries in this manual.
this manual.
tsearch, tdelete, twalk:
hsearch, hcreate, hdestroy:
records. fwtmp, wtmpfix:
frexp, ldexp, modf:
tp:
manual. man.
manprog: print entries in this
for formatting entries in this
ascii:
files. diffmk:
umask: set file-creation mode
set and get file creation
an error message file by
table. master:
information table.
regular expression compile and
eqn, neqn, checkeq: format
function.
maze: generate a

bcd: convert to antique

memcpy, memset: memory/
memset: memory/
operations. memcpy, memchr,
memcpy, memchr, memcmp,
memchr, memcmp, memset: memory/
memory.
memory. plock:
memory allocator. malloc,
memory control operations.
memory id. /remove a message
memory operations. /memchr,
memory operations.
memory segment.
memset: memory operations.
merge files.
merge or add total accounting
merge same lines of several
msg: permit or deny messages.
msgctl:
mail. mail, rmail:
mail, rmail: send mail to
mail system.
mail to.
mail to arbitrary people.
mail to users or read mail.
main memory allocator.
maintain a tags file for a C
maintain, update, and
maintainer.
maintenance commands and/
maintenance procedures.
make a delta (change) to an
make a directory.
make a directory, or a special
make a lost+found directory
make a unique file name.
make: maintain, update, and
make output single spaced.
make posters.
makekey: generate encryption
malloc, free, realloc, calloc:
man: macros for formatting
man, manprog: print entries in
manage binary search trees.
manage hash search tables.
manipulate connect accounting
manipulate parts of/
manipulate tape archive.
manprog: print entries in this
manual. man,
manual. man: macros
map of ASCII character set.
mark differences between
mask.
mask. umask:
massaging C source. /create
master device information
master: master device
match routines. regexp:
mathematical text for nroff or/
matherr: error-handling
maze.
maze: generate a maze.
media.
mem, kmem: core memory.
memcpy, memchr, memcmp,
memchr, memcmp, memcpy,
memcmp, memcpy, memset: memory
memcpy, memset: memory/
memory.
memory. plock:
memory allocator. malloc,
memory control operations.
memory id. /remove a message
memory operations. /memchr,
memory operations.
memory segment.
memset: memory operations.
merge files.
merge or add total accounting
merge same lines of several
msg: permit or deny messages.
msgctl:
mail.1
mail.1
netmail.8
netmailer.8
delivermail.8
mail.1
malloc.3c
ctags.1
make.1
ar.1
intro.1m
intro.8
delta.1
mkdir.1
mknod.2
mklost+found.1m
mktemp.3c
make.1
ssp.1
banner.1
makekey.1
malloc.3c
man.5
man.1
tsearch.3c
hsearch.3c
fwtmp.1m
frexp.3c
tp.1
man.1
man.1
man.5
ascii.5
diffmk.1
umask.1
umask.2
mkstr.1
master.4
master.4
regexp.5
eqn.1
matherr.3m
maze.6
maze.6
bcd.6
mem.7
memory.3c
memory.3c
memory.3c
memory.3c
mem.7
plock.2
malloc.3c
shmctl.2
ipcrm.1
memory.3c
shmop.2
shmget.2
memory.3c
sort.1
acctmrg.1m
paste.1
msg.1
msgctl.2

mkstr: create an error
receive: receive
send: send
msgop:
msgget: get
or shared/ ipcrm: remove a
msg: permit or deny
sys_nerr: system error
message file by massaging C/
message from a socket.
message from a socket.
message operations.
message queue.
message queue, semaphore set
messages.
messages. /errno, sys_errlist,
mkdir: make a directory.
mkfs: construct a file system.
mkfs512: construct a file
mklost+found: make a
mknod: build special file.
mknod: make a directory, or a
mkstr: create an error message
mktemp: make a unique file
MM macro package for
MM macros. /print/check
mm, osdd, checkmm: print/check
mm: the MM macro package for
mmt, mvt: typeset documents,
table.
chmod: change
umask: set file-creation
chmod: change
tset: set terminal
getty: set terminal type,
bs: a compiler/interpreter for
floating-point/ frexp, ldexp,
utime: set file access and
touch: update access and
/ckpacct, dodisk, lastlogin,
profile.
uusb:

package for formatting/
/ASCII formats suitable for
rcvhex: translates
mount:
system. mount, umount:

setmnt: establish
dismount file system.
mnttab:
mmdir:
cp, ln, mv: copy, link or
lseek:
the LP request scheduler and
formatting a permuted index.
/erand48, lrand48, srand48./
msgctl: message control
msgget: get message queue.
msgop: message operations.
multiplexing.
mv: a troff macro package for
mv: copy, link or move files.
mmdir. move a directory.
mmt, i-numbers,
definitions for eqn and
mathematical text for/ eqn,
networking facilities.
system.
mkstr.1
receive.2
send.2
msgop.2
msgget.2
ipcrm.1
msg.1
perror.3c
mkdir.1
mkfs.1m
mkfs512.1m
mklost+found.1m
mknod.1m
mknod.2
mkstr.1
mktemp.3c
mm.5
mm.1
mm.1
mm.5
mmt.1
mnttab.4
chmod.1
umask.1
chmod.2
tset.1
getty.1m
bs.1
frexp.3c
utime.2
touch.1
acctsh.1m
monitor.3c
uusb.1m
moo.6
mosd.5
hex.1
rcvhex.1
mount.2
mount.1m
mount.2
setmnt.1m
mount.1m
mnttab.4
mmdir.1m
cp.1
lseek.2
lpsched.1m
mptx.5
drand48.3c
msgctl.2
msgget.2
msgop.2
select.2
mv.5
cp.1
mmdir.1m
mmt.1
ncheck.1m
eqnchar.5
eqn.1
net.5
netmail.8
netmailer.8

uucb: monitor uucp	network.	uucb.1m	assembler and link editor	output. a.out:	a.out.4
netmail: the bnet	network mail system.	netmail.8	sprintf: print formatted	output. printf, fprintf,	printf.3s
rstat:	network statistics program	rstat.1	ssp: make	output single spaced.	ssp.1
net: introduction to	networking facilities.	net.5	/acctdusg, accton, acctwtmp:	overview of accounting and/	acct.1m
a text file.	newform: change the format of	newform.1	chown: change	owner and group of a file.	chown.2
news: print	newgrp: log in to a new group.	newgrp.1	chown, chgrp: change	owner or group.	chown.1
news items.	news items.	news.1	and expand files.	pack, pcat, unpack: compress	pack.1
news: print news items.	news.1	news.1	sadc: system activity report	package. sa1, sa2,	sa.1m
process.	nice: change priority of a	nice.2	standard buffered input/output	package. stdio:	stdio.3s
priority.	nice: run a command at low	nice.1	interprocess communication	package. stdipc: standard	stdipc.3c
nl: line numbering filter.	nl: line numbering filter.	nl.1	permuted/ mptx: the macro	package for formatting a	mptx.5
list.	nlist: get entries from name	nlist.3c	documents. mm: the MM macro	package for formatting/	mm.5
nm: print name list.	nm: print name list.	nm.1	mosd: the OSDD adapter macro	package for typesetting view	mv.5
hangups (sh only).	nohup: run a command immune to	nohup.1	graphs and/ mv: a troff macro	paginator for the Tektronix	4014.1
setjmp, longjmp:	non-local goto.	setjmp.3c	4014 terminal. 4014:	parameters. disktime -	disktime.1m
drand48, erand48, lrand48,	nrand48, mrand48, jrand48,/	drand48.3c	tune floppy disk settling time	parent process IDs. /get	getpid.2
format mathematical text for	nroff: format text.	nroff.1	process, process group, and	getopt:	getopt.1
tbl: format tables for	nroff or troff. /checkeq:	eqn.1	getopt:	passwd: change login password.	passwd.1
typesetting.	nroff or troff.	tbl.1	passwd: password file.	passwd: password file.	passwd.4
constructs. deroff: remove	nroff7: text formatting and	nroff.7.1	passwd.	passwd.	passwd.1
constructs. deroff: remove	nroff/troff, tbl, and eqn	deroff.1	passwd file.	passwd file.	passwd.4
constructs. deroff: remove	null file.	null.7	passwd file entry.	passwd file entry.	passwd.3c
constructs. deroff: remove	null: the null file.	null.7	password file.	password file.	passwd.4
constructs. deroff: remove	nulladm, prctmp, prdaily,/	acctsh.1m	password file entry.	password file entry.	passwd.3c
constructs. deroff: remove	numbering filter.	nl.1	password file entry.	password file entry.	passwd.3c
constructs. deroff: remove	numerals to English.	number.6	password/group file checkers.	password/group file checkers.	pwck.1m
constructs. deroff: remove	numeric data in a machine/	sputl.3x	paste: merge same lines of	paste: merge same lines of	paste.1
constructs. deroff: remove	object file. dump:	dump.1	path names. basename,	path names. basename,	basename.1
constructs. deroff: remove	object file.	size.1	pathname of current working	pathname of current working	getcwd.3c
constructs. deroff: remove	object files into ASCII	hex.1	fgrep: search a file for a	pattern. grep, egrep,	grep.1
constructs. deroff: remove	object library. lorder:	lorder.1	processing language. awk:	pattern scanning and	awk.1
constructs. deroff: remove	object, or other binary file.	strings.1	signal.	pause: suspend process until	pause.2
constructs. deroff: remove	octal dump.	od.1	expand files. pack,	pcat, unpack: compress and	pack.1
constructs. deroff: remove	od: octal dump.	od.1	a process. popen,	pclose: initiate pipe to/from	popen.3s
constructs. deroff: remove	only. nohup: run a command	nohup.1	value about your/ m68k,	pdpl1, u3b, vax: provide truth	machid.1
constructs. deroff: remove	on/off the extended errors in	exterr.1	msg:	permit or deny messages.	msg.1
constructs. deroff: remove	onto a remote machine..	put.1c	macro package for formatting a	permuted index. mptx: the	mptx.5
constructs. deroff: remove	onto a remote machine..	put7.1c	ptx:	permuted index.	ptx.1
constructs. deroff: remove	open a stream.	fopen.3s	format. acct:	per-process accounting file	acct.4
constructs. deroff: remove	open file descriptor.	dup.2	acctcms: command summary from	per-process accounting/	acctcms.1m
constructs. deroff: remove	open for reading or writing.	open.2	sys_nerr: system error/	perusal filter for crt	more.1
constructs. deroff: remove	open: open for reading or	open.2	viewing. more: file	phototypesetter simulator.	tc.1
constructs. deroff: remove	operating system profiler.	profiler.1m	tc:	phys: allow a process to	phys.2
constructs. deroff: remove	operation routines. /tgoto,	termcap.3	access physical addresses.	physical addresses. phys:	phys.2
constructs. deroff: remove	operations. memccpy, memchr,	memory.3c	allow a process to access	pieces.	split.1
constructs. deroff: remove	operations.	msgctl.2	split: split a file into	channel.	pipe.2
constructs. deroff: remove	operations.	msgop.2	tee:	pipe fitting.	tee.1
constructs. deroff: remove	operations.	semctl.2	popen, pclose: initiate	pipe to/from a process.	popen.3s
constructs. deroff: remove	operations.	semop.2	fish:	play "Go Fish".	fish.6
constructs. deroff: remove	operations.	shmctl.2	life:	play the game of life.	life.6
constructs. deroff: remove	operations.	shmop.2	worm:	Play the growing worm game.	worm.6
constructs. deroff: remove	operations. /strpbrk, strspn,	string.3c	data in memory.	plck: lock process, text, or	plck.2
constructs. deroff: remove	operator.	join.1	subroutines.	plot: graphics interface.	plot.4
constructs. deroff: remove	optimal access time.	dcopy.1m	images.	plot: graphics interface	plot.3x
constructs. deroff: remove	option letter from argument	getopt.3c	lseek: move read/write file	pnch: file format for card	pnch.4
constructs. deroff: remove	options.	fcntl.5	ftell: reposition a file	pointer.	lseek.2
constructs. deroff: remove	options.	getopt.1	to/from a process.	pointer in a stream. /rewind,	fseek.3s
constructs. deroff: remove	options for a terminal.	stty.1	data base of terminal types by	popen, pclose: initiate pipe	popen.3s
constructs. deroff: remove	ordering relation for an	lorder.1	basename, dirname: deliver	port. ttytype:	ttytype.4
constructs. deroff: remove	ordinary file. mknod: make	mknod.2	banner: make	portions of path names.	basename.1
constructs. deroff: remove	oriented (visual) display	vi.1	logarithm,/ exp, log, log10,	posters.	banner.1
constructs. deroff: remove	OSDD adapter macro package for	mosd.5	/sqrt: exponential, logarithm,	power, square root functions.	exp.3m
constructs. deroff: remove	osdd, checkmm: print/check	mm.1			exp.3m
constructs. deroff: remove	out-going terminal line/	dial.3c			

brc, bcheckrc, rc, powerfail: system/ brc.1m
 pr: print files. pr.1
 /lastlogin, monacct, nulladm, prctmp, prdaily, prtacct,/ acctsh.1m
 /monacct, nulladm, prctmp, prdaily, prtacct, runacct,/ acctsh.1m
 for troff. cw, checkcw: prepare constant-width text cw.1
 monitor: prepare execution profile. monitor.3c
 cpp: the C language preprocessor. cpp.1
 unget: undo a previous get of an SCCS file. unget.1
 operating/ prfld, prfstat, prfscap, prfpr: operating/ profiler.1m
 /prfstat, prfdc, prfscap, prfpr: operating system/ profiler.1m
 system/ prfld, prfstat, prfdc, prfscap, prfpr: operating profiler.1m
 graphical/ gps: graphical primitive string, format of gps.4
 types: primitive system data types. types.5
 interesting, adage. fortune: print a random, hopefully fortune.6
 prs: print an SCCS file. prs.1
 date: print and set the date. date.1
 cal: print calendar. cal.1
 of a file. sum: print checksum and block count sum.1
 editing activity. sact: print current SCCS file sact.1
 man, manprog: print entries in this manual. man.1
 cat: concatenate and print files. cat.1
 pr: print files. pr.1
 printf, fprintf, sprintf: print formatted output. printf.3s
 banner7: print large banner on printer. banner7.1
 lpstat: print LP status information. lpstat.1
 nm: print name list. nm.1
 system hostname: set or print name of current host hostname.1
 System. uname: print name of current UNIX uname.1
 news: print news items. news.1
 printenv: print out the environment. printenv.1
 file(s). acctcom: search and print process accounting acctcom.1
 pstat: print system facts. pstat.1m
 names. id: print user and group IDs and id.1
 object, or/ strings: find the printable strings in an strings.1
 formatted/ mm, osdd, checkmm: print/check documents mm.1
 environment: printenv: print out the printenv.1
 banner7: print large banner on printer. banner7.1
 requests to an LP line lp: printer. /cancel: send/cancel lp.1
 printer spooler. lpr.1
 disable: enable/disable LP printers. enable, enable.1
 print formatted output. printf, fprintf, sprintf: printf.3s
 nice: run a command at low priority. nice.1
 nice: change priority of a process. nice.2
 exit, _exit: terminate process. exit.2
 fork: create a new process. fork.2
 inittab: script for the init process. inittab.4
 kill: terminate a process. kill.1
 nice: change priority of a process. nice.2
 initiate pipe to/from a process. popen, pclose: popen.3s
 wait: await completion of process. wait.1
 errors. errpt: process a report of logged errpt.1m
 acct: enable or disable process accounting. acct.2
 acctprc1, acctprc2: process accounting. acctprc.1m
 acctcom: search and print process accounting file(s). acctcom.1
 times. times: get process and child process times.2
 init, telinit: process control/ init.1m
 timex: time a command; report process data and system/ timex.1
 /getpgrp, getppid: get process group, and parent/ getpid.2
 setpgrp: set process group ID. setpgrp.2
 process group, and parent process IDs. /get process, getpid.2
 kill: send a signal to a process or a group of/ kill.2
 getpid, getpgrp, getppid: get process, process group, and/ getpid.2

ps: report process status. ps.1
 memory. plock: lock process, text, or data in plock.2
 times: get process and child process times. times.2
 addresses. phys: allow a process to access physical phys.2
 wait: wait for child process to stop or terminate. wait.2
 ptrace: process trace. ptrace.2
 pause: suspend process until signal. pause.2
 list of file systems processed by fsck. checklist: checklist.4
 to a process or a group of processes. /send a signal kill.2
 killall: kill all active processes. killall.1m
 structure. fuser: identify processes using a file or file fuser.1m
 shutdown: terminate all processing. shutdown.1m
 awk: pattern scanning and processing language. awk.1
 m4: macro processor. m4.1
 provide truth value about your processor type. /u3b, vax: machid.1
 alarm: set a process's alarm clock. alarm.2
 profile: prof: display profile data. prof.1
 profile: execution time profil.2
 monitor: prepare execution profile. monitor.3c
 profil: execution time profile. profil.2
 prof: display profile data. prof.1
 environment at login time. profile: setting up an profle.4
 prfpr: operating system profiler. /prfdc, prfscap, profiler.1m
 sadp: disk access profiler. sadp.1
 standard/restricted command programming language. /the sh.1
 ip: Internet Protocol. ip.5
 Internet Transmission Control Protocol. tcp: tcp.5
 udp: Internet User Datagram Protocol. udp.5
 inet: Internet protocol family. inet.5
 arithmetic: provide drill in number facts. arithmetic.6
 for reading or/ lockf: provide exclusive file regions lockf.2
 m68k, pdpl1, u3b, vax: provide truth value about your/ machid.1
 true, false: provide truth values. true.1
 prs: print an SCCS file. prs.1
 /nulladm, prctmp, prdaily, prtacct, runacct, shutacct,/ acctsh.1m
 ps: report process status. ps.1
 /generate uniformly distributed pseudo-random numbers. drand48.3c
 pstat: print system facts. pstat.1m
 ptrace: process trace. ptrace.2
 ptx: permuted index. ptx.1
 stream. ungetc: push character back into input ungetc.3s
 remote machine.. put7: puts a file onto a put7.1c
 put character or word on a/ putc, putchar, fputc, putw: putc.3s
 character or word on a/ putc, putchar, fputc, putw: put putc.3s
 entry. putpwent: write password file putpwent.3c
 machine.. put: puts a file onto a remote put.1c
 machine.. put7: puts a file onto a remote put7.1c
 stream. puts, fputs: put a string on a puts.3s
 getutent, getutid, getutline, pututline, setutent, endutent,/ getut.3c
 a/ putc, putchar, fputc, putw: put character or word on putc.3s
 file checkers. pwck, grpck: password/group pwck.1m
 pwd: working directory name. pwd.1
 qsort: quicker sort. qsort.3c
 queue. msgget: get message msgget.2
 ipcrm: remove a message queue, semaphore set or shared/ ipcrm.1
 qsort: quicker sort. qsort.3c
 quiz: test your knowledge. quiz.6
 by name or address. rhost, raddr: look up internet hosts rhost.3
 display. rain: animated raindrops rain.6
 rain: animated raindrops display. rain.6
 random-number generator. rand, srand: simple rand.3c
 adage. fortune: print a random, hopefully interesting, fortune.6
 rand, srand: simple random-number generator. rand.3c
 fsplit: split fortran, ratfor, or efl files. fsplit.1

initialization/ brc, bcheckrc, rc, powerfail: system brc.1m
rcp: remote file copy rcp.1
S-records from downloading/
getpass: read a password. getpass.3c
read: read from file. read.2
rmail: send mail to users or
line: read mail. mail, mail.1
read one line. line.1
read: read from file. read.2
exclusive file regions for
open: open for reading or writing. /provide
lockf.2
lseek: move reading or writing. open.2
allocator. malloc, free, read/write file pointer. lseek.2
reboot: realloc, calloc: main memory malloc.3c
reboot the system. reboot.2
reboot the system. reboot.2
specify what to do upon receipt of a signal. signal: signal.2
receive: receive message from a socket. receive.2
receive: receive message from a socket. receive.2
from per-process accounting records. /command summary acctcms.1m
manipulate connect accounting records. fwtmp, wtmpfix: fwtmp.1m
errdead: extract error records from dump. errdead.1m
tape. frec: recover files from a backup frec.1m
ed, red: text editor. ed.1
generate C program cross reference. cxref: cxref.1
execute regular expression. regcmp, regex: compile and regcmp.3x
compile. regcmp: regular expression regcmp.1
make: maintain, update, and regenerate groups of programs. make.1
regular expression. regcmp, regex: compile and execute regcmp.3x
compile and match routines. regexp: regular expression regexp.5
lockf: provide exclusive file regions for reading or/ lockf.2
regex: compile and execute regular expression. regcmp, regcmp.3x
regcmp: regular expression compile. regcmp.1
match routines. regexp: regular expression compile and regexp.5
requests. accept, reject: allow/prevent LP accept.1m
sorted files. comm: select or reject lines common to two comm.1
lorder: find ordering relation for an object/ lorder.1
join: relational database operator. join.1
strip: remove symbols and relocation bits. strip.1
/fmod, fabs: floor, ceiling, remainder, absolute value/ floor.3m
calendar: reminder service. calendar.1
rcp: remote file copy rcp.1
rlogin: remote login rlogin.1
put: puts a file onto a remote machine.. put.1c
put7: puts a file onto a remote machine.. put7.1c
take: takes a file from a remote machine. take.1c
take7: takes a file from a remote machine.. take7.1c
remsh: remote shell remsh.1
ct: spawn getty to a remote terminal. ct.1c
file. rmdel: remove a delta from an SCCS rmdel.1
semaphore set or/ ipcrm: remove a message queue, ipcrm.1
unlink: remove directory entry. unlink.2
rm, rmdir: remove files or directories. rm.1
eqn constructs. deroff: remove nroff/troff, tbl, and deroff.1
bits. strip: remove symbols and relocation strip.1
remsh: remote shell remsh.1
check and interactive repair. /system consistency fsck.1m
uniq: report repeated lines in a file. uniq.1
clock: report CPU time used. clock.3c
communication/ ipc: report inter-process ipc.1
blocks. df: report number of free disk df.1m
errpt: report of logged errors. errpt.1m
frequencies in a file. freq: report on character freq.1
sa2, sadc: system activity report package. sal, sar.1m
timex: time a command; report process data and system/ timex.1
ps: report process status. ps.1

file. uniq: report repeated lines in a uniq.1
sar: system activity reporter. sar.1
files. version: reports version number of version.1
stream. fseek, rewind, ftell: reposition a file pointer in a fseek.3s
/lpmove: start/stop the LP request scheduler and move/ lp sched.1m
reject: allow/prevent LP requests. accept, accept.1m
LP request scheduler and move requests. /start/stop the lp sched.1m
lp, cancel: send/cancel requests to an LP line/ lp.1
to a sensible state. reset: reset the teletype bits reset.1
sensible state. reset: reset the teletype bits to a reset.1
a socket. socketaddr: return address associated with socketaddr.2
abs: return integer absolute value. abs.3c
logname: return login name of user. logname.3x
name. getenv: return value for environment getenv.3c
stat: data returned by stat system call. stat.5
configuration/ uvar: returns system-specific uvar.2
col: filter reverse line-feeds. col.1
file pointer in a/ fseek, rewind, ftell: reposition a fseek.3s
creat: create a new file or rewrite an existing one. creat.2
hosts by name or address. rhost, raddr: look up internet rhost.3
rlogin: remote login rlogin.1
directories. rm, rmdir: remove files or rm.1
read mail. mail, rmail: send mail to users or mail.1
SCCS file. rmdel: remove a delta from an rmdel.1
directories. rm, rmdir: remove files or rm.1
Escape from the automatic robots. autorobots: autorobots.6
Try to escape the killer robots. chase: chase.6
robots: Escape from the robots. robots.6
robots: robots: Escape from the robots.6
chroot: change root directory. chroot.2
chroot: change root directory for a command. chroot.1m
logarithm, power, square root functions. /exponential, exp.3m
expression compile and match routines. regexp: regular regexp.5
terminal independent operation routines. /tgoto, tputs: termcap.3
standard/restricted/ sh, rsh: shell, the sh.1
program rstat: network statistics rstat.1
nice: run a command at low priority. nice.1
hangups (sh) nohup: run a command immune to nohup.1
runacct: run daily accounting. runacct.1m
runacct: run daily accounting. runacct.1m
/prctmp, prdaily, prtacct, runacct, shutacct, startup,/ acctsh.1m
local machines. show host status of ruptime.1
local machines rwho: who is logged in on rwho.1
activity report package. sal, sa2, sadc: system sar.1m
report package. sal, sa2, sadc: system activity sar.1m
editing activity. sact: print current SCCS file sact.1
package. sal, sa2, sadc: system activity report sar.1m
sadb: disk access profiler. sadp.1
sag: system activity graph. sag.1g
sar: system activity reporter. sar.1
space allocation. brk, sbrk: change data segment brk.2
formatted input. scanf, fscanf, sscanf: convert scanf.3s
bfs: big file scanner. bfs.1
language. awk: pattern scanning and processing awk.1
the delta commentary of an SCCS delta. cdc: change cdc.1
comb: combine SCCS deltas. comb.1
make a delta (change) to an SCCS file. delta: delta.1
get: get a version of an SCCS file. get.1
prs: print an SCCS file. prs.1
rmdel: remove a delta from an SCCS file. rmdel.1
compare two versions of an SCCS file. sccsdiff: sccsdiff.1
sccsfile: format of SCCS file. sccsfile.4
undo a previous get of an SCCS file. unget: unget.1
val: validate SCCS file. val.1

sact: print current
 admin: create and administer
 what: identify
 of an SCCS file.
 /start/stop the LP request
 clear: clear terminal
 twinkle: twinkle stars on the
 terminals. se:
 display editor/ vi, view:
 inittab:
 system initialization shell
 program.
 terminals.
 bsearch: binary
 grep, egrep, fgrep:
 accounting file(s). acctcom:
 lsearch: linear
 hcreate, hdestroy: manage hash
 tdelete, twalk: manage binary
 /mrand48, jrand48, srand48,
 shmget: get shared memory
 brk, sbrk: change data
 to two sorted files. comm:
 multiplexing.
 greek:
 of a file. cut: cut out
 file. dump: dump
 semctl:
 semop:
 ipcrm: remove a message queue,
 semget: get set of
 operations.
 a group of processes. kill:
 mail. mail, rmail:
 send:
 socket.
 line printer. lp, cancel:
 reset the teletype bits to a
 stream.
 IDs. setuid,
 getgrent, getgrgid, getgrnam,
 cpu.
 goto.
 encryption. crypt,
 getpwent, getpwuid, getpwnam,
 login time. profile:
 gettydefs: speed and terminal
 disktime - tune floppy disk
 group IDs.
 /getutid, getutline, pututline,
 data in a machine/ sputl,
 standard/restricted command/
 operations. shmctl:
 queue, semaphore set or
 shmop:
 shmget: get
 remsh: remote
 system: issue a
 SCCS file editing activity. sact.1
 SCCS files. admin.1
 SCCS files. what.1
 sccsdiff: compare two versions sccsdiff.1
 sccsfile: format of SCCS file. sccsfile.4
 scheduler and move requests. lpsched.1m
 screen. clear.1
 screen. twinkle.6
 screen editor for video se.1
 screen oriented (visual) vi.1
 script for the init process. inittab.4
 scripts. /rc, powerfail: brc.1m
 sdiff: side-by-side difference sdiff.1
 se: screen editor for video se.1
 search. bsearch.3c
 search a file for a pattern. grep.1
 search and print process acctcom.1
 search and update. lsearch.3c
 search tables. hsearch, hsearch.3c
 search trees. tsearch, tsearch.3c
 sed: stream editor. sed.1
 seed48, lcong48: generate/
 segment. drand48.3c
 shmget.2
 segment space allocation. brk.2
 select or reject lines common comm.1
 select: synchronous i/o select.2
 select terminal filter. greek.1
 selected fields of each line cut.1
 selected parts of an object dump.1
 semaphore control operations. semctl.2
 semaphore operations. semop.2
 semaphore set or shared memory/
 semaphores. ipcrm.1
 semget.2
 semctl: semaphore control semctl.2
 semget: get set of semaphores. semget.2
 semop: semaphore operations. semop.2
 send a signal to a process or kill.2
 send mail to users or read mail.1
 send message from a socket. send.2
 send: send message from a send.2
 send/cancel requests to an LP lp.1
 sensible state. reset: reset.1
 setbuf: assign buffering to a setbuf.3s
 setgid: set user and group setuid.2
 setgrent, endgrent: get group/
 sethostname: set name of host sethostname.2
 setjmp, longjmp: non-local setjmp.3c
 setkey, encrypt: generate DES crypt.3c
 setmnt: establish mount table. setmnt.1m
 setpgrp: set process group ID. setpgrp.2
 setpwent, endpwent: get/ getpwent.3c
 setting up an environment at profile.4
 settings used by getty. gettydefs.4
 settling time parameters. disktime.1m
 setuid, setgid: set user and setuid.2
 setutent, endutent, utmpname:/ getut.3c
 sgetl: access long numeric sputl.3x
 sh, rsh: shell, the sh.1
 shared memory control shmctl.2
 shared memory id. /a message ipcrm.1
 shared memory operations. shmop.2
 shared memory segment. shmget.2
 shell remsh.1
 shell command. system.3s

with C-like syntax. csh: a
 shutacct, startup, turnacct:
 system initialization
 command programming/ sh, rsh:
 operations.
 segment.
 operations.
 /prdaily; prtacct, runacct,
 processing.
 program. sdiff:
 login:
 sign on. login.1
 signal. pause.2
 what to do upon receipt of
 upon receipt of a signal.
 of processes. kill: send a
 ssignal, gsignal: software
 lex: generate programs for
 generator. rand, srand:
 tc: phototypesetter
 atan, atan2: trigonometric/
 ssp: make output
 functions.
 size:
 an interval.
 interval.
 documents, view graphs, and
 typesetting view graphs and
 current/ ttyslot: find the
 spline: interpolate
 sno:
 SNOBOL interpreter. sno.1
 socket. accept: accept.2
 socket. connect: connect.2
 socket. receive: receive.2
 send: send message from a
 address associated with a
 communication.
 associated with a socket.
 loop:
 ssignal, gsignal:
 qsort: quicker
 tsort: topological
 sort:
 sort and/or merge files. sort.1
 sort: sort and/or merge files. sort.1
 sorted files. comm: select comm.1
 source. /create an error mkstr.1
 space allocation. brk.2
 ssp: make output single
 terminal. ct:
 fspec: format
 the extended errors in the
 receipt of a signal. signal:
 /set terminal type, modes,
 used by getty. gettydefs:
 hashcheck: find spelling/
 spelling/ spell, hashmake,
 spellin, hashcheck: find
 spelling errors. /hashmake,
 spline: interpolate smooth spline.1g
 split. csplit.1
 split a file into pieces. split.1
 efl files. fsplit:
 pieces.
 uuclean: uucp
 shell (command interpreter) csh.1
 shell procedures for/ /runacct, acctsh.1m
 shell scripts. /rc, powerfail: brc.1m
 shell, the standard/restricted sh.1
 shmctl: shared memory control shmctl.2
 shmget: get shared memory shmget.2
 shmop: shared memory shmop.2
 shutacct, startup, turnacct:/ acctsh.1m
 shutdown: terminate all shutdown.1m
 side-by-side difference sdiff.1
 sign on. login.1
 signal. pause.2
 signal. signal: specify signal.2
 signal: specify what to do signal.2
 signal to a process or a group kill.2
 signals. ssignal.3c
 simple lexical tasks. lex.1
 simple random-number rand.3c
 simulator. tc.1
 sin, cos, tan, asin, acos, trig.3m
 single spaced. ssp.1
 sinh, cosh, tanh: hyperbolic sinh.3m
 size of an object file. size.1
 size: size of an object file. size.1
 sleep: suspend execution for sleep.1
 sleep: suspend execution for sleep.3c
 slides. mmt, mvt: typeset mmt.1
 slides. /macro package for mv.5
 slot in the utmp file of the ttyslot.3c
 smooth curve. spline.1g
 sno: SNOBOL interpreter. sno.1
 SNOBOL interpreter. sno.1
 socket. accept: accept.2
 socket. connect: connect.2
 socket. receive: receive.2
 socket. send.2
 socket. socketaddr: return socketaddr.2
 socket: create an endpoint for socket.2
 socketaddr: return address socketaddr.2
 software loopback interface. lo.5
 software signals. ssignal.3c
 sort. qsort.3c
 sort. tsort.1
 sort and/or merge files. sort.1
 sort: sort and/or merge files. sort.1
 sorted files. comm: select comm.1
 source. /create an error mkstr.1
 space allocation. brk.2
 ssp: make output single
 terminal. ct:
 fspec: format
 specification in text files. fspec.4
 specified device. /turn on/off exterr.1
 specify what to do upon signal.2
 speed, and line discipline. getty.1m
 speed and terminal settings gettydefs.4
 spell, hashmake, spellin, spell.1
 spellin, hashcheck: find spell.1
 spelling errors. /hashmake, spell.1
 spline: interpolate smooth spline.1g
 split. csplit.1
 split a file into pieces. split.1
 split fortran, ratfor, or fsplit.1
 split: split a file into split.1
 spool directory clean-up. uuclean.1m

lpr: line printer
 lpadmin: configure the LP output. printf, fprintf, numeric data in a machine/power,/ exp, log, log10, pow, exponential, logarithm, power, generator. rand, /nrand48, mrand48, jrand48, formats suitable for Motorola rcvhex: translates Motorola input. scanf, fscanf, signals. spaced.
 package. stdio: standard buffered input/output communication/ stdipc: standard interprocess sh, rsh: shell, the twinkle: twinkle
 lpsched, lpshut, lpmove: boot: /prtacct, runacct, shutacct, system call.
 stat: data returned by ustat: get file system ff: list file names and rstat: network communication facilities ps: report process stat, fstat: get file lpstat: print LP feof, clearerr, fileno: stream control. uustat: uucp ruptime: show host input/output package. communication package.
 wait for child process to strncmp, strcpy, strncpy,/ /strcpy, strncpy, strlen, strncpy,/ strcat, strncat, /strncat, strcmp, strncmp, /strchr, strpbrk, strspn, fflush: close or flush a fopen, freopen, fdopen: open a reposition a file pointer in a get character or word from fgets: get a string from a put character or word on a puts, fputs: put a string on a setbuf: assign buffering to a push character back into input sed: /feof, clearerr, fileno: long integer and base-64 ASCII convert date and time to floating-point number to gps: graphical primitive gets, fgets: get a puts, fputs: put a strspn, strcspn, strtok: string operations. /strpbrk, string to floating-point atof: convert ASCII strtol, atol, atoi: convert strings in an object, or strings: find the printable

spooler. lpr.1
 spooling system. lpadmin.1m
 sprintf: print formatted printf.3s
 sputl, sgetl: access long sputl.3x
 sqrt: exponential, logarithm, exp.3m
 square root functions. /sqrt: exp.3m
 srand: simple random-number rand.3c
 srand48, seed48, lcong48:/ drand48.3c
 S-record downloading. /ASCII hex.1
 S-records from downloading/ rcvhex.1
 sscanf: convert formatted scanf.3s
 ssignal, gsignal: software ssignal.3c
 ssp: make output single ssp.1
 standard buffered input/output stdio.3s
 standard interprocess stdipc.3c
 standard/restricted command/ sh.1
 stars on the screen. twinkle.6
 start/stop the LP request/ lpsched.1m
 startup procedures. boot.8
 startup, turnacct: shell/ acctsh.1m
 stat: data returned by stat stat.5
 stat, fstat: get file status. stat.2
 stat system call. stat.5
 statistics. ustat.2
 statistics for a file system. ff.1m
 statistics program rstat.1
 status. /report inter-process ipc.1
 status. ps.1
 status. stat.2
 status information. lpstat.1
 status inquiries. ferror. ferror.3s
 status inquiry and job uustat.1c
 status of local machines ruptime.1
 stdio: standard buffered stdio.3s
 stdipc: standard interprocess stdipc.3c
 stime: set time. stime.2
 stop or terminate. wait: wait.2
 strcat, strncat, strcmp, string.3c
 strchr, strchr, strpbrk,/ string.3c
 strcmp, strncmp, strcpy, string.3c
 strcpy, strncpy, strlen,/ string.3c
 strcspn, strtok: string/ string.3c
 stream. fclose, fclose.3s
 stream. fopen.3s
 stream. fseek, rewind, ftell: fseek.3s
 stream. /getchar, fgetc, getw:getc.3s
 stream. gets, gets.3s
 stream. /putchar, fputc, putw:putc.3s
 stream. puts.3s
 stream. setbuf.3s
 stream. ungetc: ungetc.3s
 stream editor. sed.1
 stream status inquiries. ferror.3s
 string. /l64a: convert between a64l.3c
 string. /asctime, tzset: ctime.3c
 string. /fcvt, gcvt: convert ecvt.3c
 string, format of graphical/ gpts.4
 string from a stream. gets.3s
 string on a stream. puts.3s
 string operations. /strpbrk, string.3c
 string to floating-point atof.3c
 string to integer. strtol.3c
 strings: find the printable strings.1
 strings in an object, or other/ strings.1

relocation bits. strip.1
 /strncmp, strcpy, strncpy, string.3c
 strcpy, strncpy,/ strcat, string.3c
 strcat, strncat, strcmp, string.3c
 strncmp, strcpy, strncpy,/ string.3c
 /strcmp, strncmp, strcpy, string.3c
 /strlen, strchr, strchr, string.3c
 /strcpy, strlen, strchr, string.3c
 /strchr, strchr, strpbrk, string.3c
 /strpbrk, strspn, strcspn, string.3c
 string to integer. strtol.3c
 processes using a file or file fuser.1m
 terminal. stty.1
 stty: set the options for a su.1
 another user. su.1
 plot: graphics interface plot.3x
 intro: introduction to intro.3
 subroutines and libraries. paste.1
 subsequent lines of one file. hex.1
 suitable for Motorola S-record/ sum7.1
 file. sum7: sum and count blocks in a sumdir.1
 the files in the/ sumdir: sum and count characters in sum.1
 count of a file. sum: print checksum and block sum7.1
 a file. sum7: sum and count blocks in sumdir.1
 characters in the files in/ du.1
 du: summarize disk usage. acctcms.1m
 accounting/ acctcms: command sync.1
 sync: update the sync.2
 sync: update su.1
 su: become sleep.1
 interval. sleep: suspend execution for an sleep.3c
 interval. sleep: suspend execution for pause.2
 pause: suspend process until signal. swab.3c
 swab: swap bytes. swab.3c
 swap bytes. strip.1
 strip: remove sync.2
 sync: update super-block. sync.1
 sync: update the super block. select.2
 select: synchronous i/o multiplexing. csh.1
 interpreter) with C-like perror.3c
 error/ perror, errno, perror.3c
 perror, errno, sys_errlist, uvar.2
 information. uvar: returns uuto.1c
 uuto, uupick: public UNIX master.4
 master device information mnttab.4
 mnttab: mounted file system table.1m
 setmnt: establish mount hosts.7
 /etc/hosts: host hsearch.3c
 hdestroy: manage hash search tbl.1
 tbl: format tabs.1
 tabs: set tabs.1
 tabs: set tabs on a terminal. ctags.1
 ctags: maintain a tail.1
 a file. take.1c
 remote machine. take: takes a file from a take7.1c
 remote machine.. take7: takes a file from a take.1c
 machine. take: takes a file from a remote take7.1c
 machine.. take7: takes a file from a remote trig.3m
 trigonometric/ sin, cos, sinh.3m
 sinh, cosh, frec.1m
 tanh: hyperbolic functions. tp.1
 recover files from a backup hpio.1
 tp: manipulate tar.1
 hpio: HP 2645A terminal tp.4
 tar: tape format. filesave.1m
 tp: magnetic tar.1
 file system backup. filesave, filesave.1m
 tapesave: daily/weekly UNIX tar.1
 tar: tape file archiver. tar.1

programs for simple lexical
 deroff: remove nroff/troff,
 or troff.

Control Protocol
 search trees. tsearch,

4014: paginator for the
 state. reset: reset the
 last logins of users and
 initialization. init,
 temporary file. tmpnam,
 tmpfile: create a
 tmpnam: create a name for a
 terminals.
 data base.
 for the Tektronix 4014
 functions of the DASI 450
 ct: spawn getty to a remote
 generate file name for
 stty: set the options for a
 tabs: set tabs on a
 isatty: find name of a
 animate worms on a display
 termcap:
 greek: select
 /tgetstr, tgoto, tputs:
 termio: general
 tty: controlling
 dial: establish an out-going
 tset: set
 clear: clear
 getty. gettydefs: speed and
 hpio: HP 2645A
 and line/ getty: set
 ttytype: data base of
 functions of DASI 300 and 300s
 of HP 2640 and 2621-series
 se: screen editor for video
 term: conventional names for
 tty: get the
 for child process to stop or
 kill:
 shutdown:
 exit, _exit:
 daemon. errstop:
 interface.
 command.
 quiz:
 nroff: format
 troff: typeset
 ed, red:
 ex, edit:
 change the format of a
 fspec: format specification in
 /checkeq: format mathematical
 prepare constant-width
 typesetting. nroff7:
 typesetting. troff7:
 plock: lock process,
 tgetstr, tgoto, tputs:/
 tputs:/ tgetent, tgetnum,
 tgoto, tputs:/ tgetent,
 tgetent, tgetnum, tgetflag,

tasks. lex: generate lex.1
 tbl, and eqn constructs. deroff.1
 tbl: format tables for nroff tbl.1
 tc: phototypesetter simulator. tc.1
 tcp: Internet Transmission tcp.5
 tdelete, twalk: manage binary
 tee: pipe fitting. tee.1
 Tektronix 4014 terminal. 4014.1
 teletype bits to a sensible reset.1
 teletypes. last: indicate last.1
 telinit: process control init.1m
 tmpnam: create a name for a tmpnam.3s
 temporary file. tmpfile.3s
 temporary file. tmpnam, tmpnam.3s
 term: conventional names for term.5
 termcap: terminal capability termcap.5
 terminal. 4014: paginator 4014.1
 terminal. 450: handle special 450.1
 terminal. ct.1c
 terminal. ctermid: ctermid.3s
 terminal. stty.1
 terminal. tabs.1
 terminal. ttyname, ttyname.3c
 terminal. worms: worms.6
 terminal capability data base. termcap.5
 terminal filter. greek.1
 terminal independent operation/ termcap.3
 terminal interface. termio.7
 terminal interface. tty.7
 terminal line connection. dial.3c
 terminal modes. tset.1
 terminal screen. clear.1
 terminal settings used by gettydefs.4
 terminal tape file archiver. hpio.1
 terminal type, modes, speed, getty.1m
 terminal types by port. ttytype.4
 terminals. /handle special 300.1
 terminals. /special functions hp.1
 terminals. se.1
 terminals. term.5
 terminal's name. tty.1
 terminate. wait: wait wait.2
 terminate a process. kill.1
 terminate all processing. shutdown.1m
 terminate process. exit.2
 terminate the error-logging errstop.1m
 termio: general terminal termio.7
 test: condition evaluation test.1
 test your knowledge. quiz.6
 text. nroff.1
 text. troff.1
 text editor. ed.1
 text editor ex.1
 text file. newform: newform.1
 text files. fspec.4
 text for nroff or troff. eqn.1
 text for troff. cw, checkcw: cw.1
 text formatting and nroff7.1
 text formatting and troff7.1
 text, or data in memory. plock.2
 tgetent, tgetnum, tgetflag, termcap.3
 tgetflag, tgetstr, tgoto, termcap.3
 tgetnum, tgetflag, tgetstr, termcap.3
 tgetstr, tgoto, tputs:/ termcap.3

/tgetnum, tgetflag, tgetstr, tgoto, tputs: terminal/ termcap.3
 ttt, cubic: tic-tac-toe. ttt.6
 time. at: at.1
 execute commands at a later
 systems for optimal access
 up an environment at login
 time. dcopy: copy file dcopy.1m
 time. profile: setting profile.4
 time. stime.2
 time: get
 time: set
 time: get
 time: time a command. time.1
 data and system/ timex:
 time a command; report process timex.1
 time: get time. time.2
 - tune floppy disk settling
 profil: execution
 time parameters. disktime disktime.1m
 time profile. profil.2
 time: time a command. time.1
 time to string. /asctime, time.3c
 time used. clock.3c
 get process and child process
 file access and modification
 process times.
 update access and modification
 process data and system/
 file.
 for a temporary file.
 /tolower, _toupper, _tolower,
 popen, pclose: initiate pipe
 toupper, tolower, _toupper,
 toascii: translate/ _toupper,
 toascii: translate/ _toupper,
 tsort:
 acctmerg: merge or add
 modification times of a file.
 translate/ _toupper, tolower,
 _tolower, toascii: translate/
 acctmerg: merge or add
 modification times of a file.
 translate/ _toupper, tolower,
 _tolower, toascii: translate/
 /tgetflag, tgetstr, tgoto,
 ptrace: process
 bit, bit512: block
 /_toupper, _tolower, toascii:
 tr:
 from downloading into/ rcvhex:
 ASCII formats suitable/ hex:
 tcp: Internet
 ftw: walk a file
 twalk: manage binary search
 trek:
 tan, asin, acos, atan, atan2:
 constant-width text for
 mathematical text for nroff or
 format tables for nroff or
 typesetting view graphs/ mv: a
 typesetting.
 values.
 m68k, pdp11, u3b, vax: provide
 true, false: provide
 robots. chase:
 manage binary search trees.
 tgoto, tputs: terminal/
 tic-tac-toe. ttt.6
 at: at.1
 dcopy: copy file dcopy.1m
 profile: setting profile.4
 stime: stime.2
 time: time.2
 time a command. time.1
 time a command; report process timex.1
 time: get time. time.2
 time parameters. disktime disktime.1m
 time profile. profil.2
 time: time a command. time.1
 time to string. /asctime, time.3c
 time used. clock.3c
 times. times: times.2
 times. utime: set utime.2
 times: get process and child times.2
 times of a file. touch: touch.1
 timex: time a command; report timex.1
 tmpfile: create a temporary tmpfile.3s
 tmpnam, tmpnam: create a name tmpnam.3s
 toascii: translate characters. conv.3c
 to/from a process. popen.3s
 _tolower, toascii: translate/
 tolower, _toupper, _tolower, conv.3c
 topological sort. tsort.1
 total accounting files. acctmerg.1m
 touch: update access and touch.1
 _toupper, _tolower, toascii: conv.3c
 toupper, tolower, _toupper, conv.3c
 tp: magnetic tape format. tp.4
 tp: manipulate tape archive. tp.1
 tplot: graphics filters. tplot.1g
 tputs: terminal independent/
 tr: translate characters. tr.1
 trace. ptrace.2
 transfer data. bit.3c
 translate characters. conv.3c
 translate characters. tr.1
 translates Motorola S-records rcvhex.1
 translates object files into hex.1
 Transmission Control Protocol. tcp.5
 tree. ftw.3c
 trees. tsearch, tdelete, tsearch.3c
 trek: trekkie game. trek.6
 trekkie game. trek.6
 trigonometric functions. /cos, trig.3m
 troff. cw, checkcw: prepare cw.1
 troff. /neqn, checkeq: format eqn.1
 troff. tbl: tbl.1
 troff macro package for mv.5
 troff: typeset text. troff.1
 troff7: text formatting and troff7.1
 true, false: provide truth true.1
 truth value about your/ machid.1
 truth values. true.1
 Try to escape the killer chase.6
 tsearch, tdelete, twalk: tsearch.3c
 tset: set terminal modes. tset.1
 tsort: topological sort. tsort.1
 ttt, cubic: tic-tac-toe. ttt.6
 tty: controlling terminal tty.7
 tty: get the terminal's name. tty.1

graphics for the extended a terminal.
 utmp file of the current/ types by port.
 parameters. disktime - /runacct, shutacct, startup, trees. tsearch, tdelete, twinkle: twinkle stars on the screen.
 file: determine file value about your processor
 getty: set terminal for the extended TTY-37 types: primitive system data ttytype: data base of terminal types.
 graphs, and slides. mmt, mvt: troff:
 nroff7: text formatting and troff7: text formatting and mv: a troff macro package for /localtime, gmtime, asctime, about your/ m68k, pdp11, Protocol.
 getpw: get name from limits.
 creation mask.
 mask.
 file system. mount, UNIX system.
 UNIX System.
 ul: do file. unget: an SCCS file. into input stream. /seed48, lcong48: generate a file.
 mktemp: make a unlink system calls. link, entry.
 unlink: exercise link and umount: files. pack, pcat, lsearch: linear search and times of a file. touch: of programs, make: maintain, badblk: program to set or machines. updater: machines. updater: sync: update super-block. sync: update the super block. updater: update files between two machines.
 two machines. du: summarize disk character login name of the logname: return login name of become super-user or another the utmp file of the current write: write to another setuid, setgid: set id: print TTY-37 type-box. greek: greek.5
 ttyname, isatty: find name of ttyname.3c
 ttyslot: find the slot in the ttyslot.3c
 ttytype: data base of terminal ttytype.4
 tune floppy disk settling time disktime.1m
 turnacct: shell procedures for/ acctsh.1m
 twalk: manage binary search tsearch.3c
 twinkle stars on the screen. twinkle.6
 twinkle: twinkle stars on the twinkle.6
 type. file.1
 type. /u3b, vax: provide truth machid.1
 type, modes, speed, and line/ getty.1m
 type-box. greek: graphics greek.5
 types. types.5
 types by port. ttytype.4
 types: primitive system data types.5
 typeset documents, view mmt.1
 typeset text. troff.1
 typesetting. nroff7.1
 typesetting. troff7.1
 typesetting view graphs and/ mv.5
 tzset: convert date and time/ ctime.3c
 u3b, vax: provide truth value machid.1
 udp: Internet User Datagram udp.5
 UID. getpw.3c
 ul: do underlining. ul.1
 ulimit: get and set user ulimit.2
 umask: set and get file umask.2
 umask: set file-creation mode umask.1
 umount: mount and dismount mount.1m
 umount: unmount a file system. umount.2
 uname: get name of current uname.2
 uname: print name of current uname.1
 underlining. ul.1
 undo a previous get of an SCCS unget.1
 unget: undo a previous get of unget.1
 ungetc: push character back ungetc.3s
 uniformly distributed/ drand48.3c
 uniq: report repeated lines in uniq.1
 unique file name. mktemp.3c
 units: conversion program. units.1
 unlink: exercise link and link.1m
 unlink: remove directory unlink.2
 unlink system calls. link, link.1m
 unmount a file system. umount.2
 unpack: compress and expand pack.1
 update. lsearch.3c
 update access and modification touch.1
 update, and regenerate groups make.1
 update bad block information. badblk.1m
 update files between two updater.1
 update files between two updater.1m
 update super-block. sync.2
 update the super block. sync.1
 updater: update files between updater.1
 updater: update files between updater.1m
 usage. du.1
 user. cuserid: get cuserid.3s
 user. logname.3x
 user. su: su.1
 user. /find the slot in ttyslot.3c
 user. write.1
 user and group IDs. setuid.2
 user and group IDs and names. id.1

udp: Internet User Datagram Protocol. udp.5
 /getgid, getegid: get real user, effective user, real/ getuid.2
 environ: user environment. environ.4
 environ: user environment. environ.5
 ulimit: get and set user limits. ulimit.2
 /get real user, effective user, real group, and/ getuid.2
 wall: write to all users. wall.1m
 last: indicate last logins of users and teletypes. last.1
 mail, rmail: send mail to users or read mail. mail.1
 fuser: identify processes using a file or file/ fuser.1m
 statistics. ustat.2
 modification times. utime.2
 utmp, wtmp: utmp and wtmp entry formats. utmp.4
 utmp file entry. /setutent, getut.3c
 utmp file of the current user. ttyslot.3c
 utmp, wtmp: utmp and wtmp utmp.4
 utmpname: access utmp file/ getut.3c
 clean-up. uuclean: uucp spool directory uuclean.1m
 uucp network. uusub.1m
 uuclean: uucp spool directory clean-up. uuclean.1m
 control. uustat: uucp status inquiry and job uustat.1c
 unix copy. uucp, uulog, uuname: unix to unix uucp.1c
 copy. uucp, uulog, uuname: unix to unix copy. uucp.1c
 uupick: public UNIX uuto.1c
 uustat: uucp status inquiry uustat.1c
 uusub: monitor uucp network. uusub.1m
 uuto, uupick: public UNIX uuto.1c
 execution. uux: unix to unix command uux.1c
 configuration information. uvar: returns system-specific uvar.2
 val: validate SCCS file. val.1
 validate SCCS file. val.1
 value. abs.3c
 value about your processor/ machid.1
 value for environment name. getenv.3c
 ceiling, remainder, absolute value functions. /fabs: floor, floor.3m
 true, false: provide truth values. true.1
 vax: provide truth value about machid.1
 vc: version control. vc.1
 vchk: version checkup. vchk.1m
 vector. getopt: get getopt.3c
 verify program assertion. assert.3x
 of directory (Berkeley version). ls7: list contents ls7.1
 vchk: version checkup. vchk.1m
 vc: version control. vc.1
 version: reports version number of files. version.1
 get: get a version of an SCCS file. get.1
 number of files. version: reports version version.1
 versions of an SCCS file. sccsdiff.1
 (visual) display editor based/ vi, view: screen oriented vi.1
 se: screen editor for video terminals. se.1
 mmt, mvt: typeset documents, view graphs, and slides. mmt.1
 macro package for typesetting view graphs and slides. /troff mv.5
 display editor based on/ vi, view: screen oriented (visual) vi.1
 file perusal filter for crt viewing. more: more.1
 on/ vi, view: screen oriented (visual) display editor based vi.1
 systems with label checking. volcopy, labelit: copy file volcopy.1m
 file system: format of system volume. fs.4
 process. wait: await completion of wait.1
 or terminate. wait: wait for child process to stop wait.2
 to stop or terminate. wait: wait for child process wait.2
 ftw: walk a file tree. ftw.3c
 wall: write to all users. wall.1m
 wc: word count. wc.1

see: see	what a file has in it.	see.1
	what: identify SCCS files.	what.1
signal: signal: specify	what to do upon receipt of a	signal.2
crashes: crash:	what to do when the system	crash.8
whodo:	who is doing what.	whodo.1m
machines rwho:	who is logged in on local	rwho.1
who:	who is on the system.	who.1
	who: who is on the system.	who.1
	whodo: who is doing what.	whodo.1m
cd: change	working directory.	cd.1
chdir: change	working directory.	chdir.2
get pathname of current	working directory. getcwd:	getcwd.3c
pwd:	working directory name.	pwd.1
worm: Play the growing	worm game.	worm.6
game.	worm: Play the growing worm	worm.6
display terminal.	worms: animate worms on a	worms.6
worms: animate	worms on a display terminal.	worms.6
write:	write on a file.	write.2
putpwent:	write password file entry.	putpwent.3c
wall:	write to all users.	wall.1m
write:	write to another user.	write.1
	write; write on a file.	write.2
	write: write to another user.	write.1
file regions for reading or	writing. /provide exclusive	lockf.2
open: open for reading or	writing.	open.2
utmp, wtmp: utmp and	wtmp entry formats.	utmp.4
formats. utmp,	wtmp: utmp and wtmp entry	utmp.4
accounting records. fwtmp,	wtmpfix: manipulate connect	fwtmp.1m
hunt-the-wumpus.	wump: the game of	wump.6
list(s) and execute command.	xargs: construct argument	xargs.1
j0, j1, jn,	y0, y1, yn: Bessel functions.	bessel.3m
j0, j1, jn, y0,	y1, yn: Bessel functions.	bessel.3m
compiler-compiler.	yacc: yet another	yacc.1
j0, j1, jn, y0, y1,	yn: Bessel functions.	bessel.3m

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.

There is a table of messages associated with each error, and a routine for printing the message; see *perror*(3). All of the possible error numbers are not listed in each system call description because many errors are possible for most of the calls. 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. 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).
- 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 Mount device busy
An attempt 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.
- 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 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's table of open files 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.
- 25 ENOTTY Not a typewriter
The file mentioned in *stty* or *gtty* is not a terminal or one of the other devices to which these calls apply.
- 26 ETXTBSY Text file busy
An attempt to execute a pure-procedure program which is currently open for writing (or reading). 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.
- 29 ESPIPE Illegal seek
An *lseek* was issued to a pipe. This error should also be issued for other non-seekable devices.
- 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)).

- 55 EWOULDBLOCK Operation would block
An operation which would cause a process to block was attempted on an object in non-blocking mode (see *ioctlnew(2)*).
- 56 EINPROGRESS Operation now in progress
An operation which takes a long time to complete (such as a *connect(2)*) was started on a non-blocking object (see *ioctlnew(2)*).
- 57 EALREADY Operation already in progress
An operation was attempted on a non-blocking object which already had an operation in progress.
- 58 ENOTSOCK Socket operation on non-socket
Self-explanatory.
- 59 EDESTADDRREQ Destination address required
A required address was omitted from an operation on a socket.
- 60 EMSGSIZE Message too long
A message sent on a socket was larger than the internal message buffer.
- 61 EPROTOTYPE Protocol wrong type for socket
A protocol was specified which does not support the semantics of the socket type requested. For example, you cannot use the internet UDP protocol with type `SOCK_STREAM`.
- 62 ENOPROTOPT Protocol not available
In this incarnation of the system.
- 63 EPROTONOSUPPORT Protocol not supported
In this incarnation of the system.
- 64 ESOCKTNOSUPPORT Socket type not supported
In this incarnation of the system.
- 65 EOPNOTSUPP Operation not supported on socket
For example, trying to *accept* a connection on a datagram socket.
- 66 EPFNOSUPPORT Protocol family not supported
In this incarnation of the system.
- 67 EAFNOSUPPORT Address family not supported by protocol family
An address incompatible with the requested protocol was used. For example, you shouldn't necessarily expect to be able to use PUP Internet addresses with ARPA Internet protocols.
- 68 EADDRINUSE Address already in use
Only one usage of each address is normally permitted.
- 69 EADDRNOTAVAIL Can't assign requested address
Normally results from an attempt to create a socket with an address not on this machine.
- 70 ENETDOWN Network is down
A socket operation encountered a dead network.
- 71 ENETUNREACH Network is unreachable
A socket operation was attempted to an unreachable network.
- 72 ENETRESET Network dropped connection on reset
The host you were connected to crashed and rebooted.

- 73 ECONNABORTED Software caused connection abort
A connection abort was caused internal to your host machine.
- 74 ECONNRESET Connection reset by peer
- 55 ENOBUFS No buffer space available
For a socket or a pipe in the buffer pool.
- 76 EISCONN Socket is already connected
- 77 ENOTCONN Socket is not connected
- 78 ESHUTDOWN Can't send after socket shutdown
- 79 *unused*
- 80 ETIMEDOUT Connection timed out
Due to failure to initiate properly or because keep-alives failed.
- 81 ECONNREFUSED Connection refused
No connection could be made because the target machine actively refused it.
- 82 ELOOP Too many levels of symbolic links
A path name lookup involved more than 8 symbolic links.
- 83 ENAMETOOLONG File name too long
A component of a path name exceeded 14 characters, or an entire path name exceeded 1023 characters.
- 84 EHOSTDOWN Host is down
A socket operation encountered a defunct host.
- 85 EHOSTUNREACH No route to host
A socket operation was attempted to an unreachable host.
- 100 EDEADLOCK Locking Deadlock
Returned by *lockf(2)* system call if deadlock would occur or when locktable overflows.

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 0 to 30,000.

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 process 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 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 excluding \0 (null) and the ASCII code for / (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 non-existent 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. A process's root directory 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 is true:

The process's effective user ID is super-user.

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

The process's effective user ID does not match the user ID of the owner of the file, and the process's effective group ID does not match the group ID of the file, and 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. 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 a *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 is true:

The process's effective user ID is super-user.

The process's effective user ID matches **msg_perm.lcluid** in the data structure associated with *msqid* and the appropriate bit of the "user" portion (0600) of **msg_perm.mode** is set.

The process's effective user ID does not match **msg_perm.lcluid** and the process's effective group ID matches **msg_perm.lclgid** and the appropriate bit of the "group" portion (060) of **msg_perm.mode** is set.

The process's effective user ID does not match **msg_perm.lcluid** and the process's effective group ID does not match **msg_perm.lclgid** 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 a *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 is true:

The process's effective user ID is super-user.

The process's effective user ID matches **sem_perm.lcluid** in the data structure associated with *semid* and the appropriate bit of the "user" portion (0600) of **sem_perm.mode** is set.

The process's effective user ID does not match **sem_perm.lcluid** and the process's effective group ID matches **sem_perm.lclgid** and the appropriate bit of the "group" portion (060) of **sem_perm.mode** is set.

The process's effective user ID does not match **sem_perm.lcluid** and the process's effective group ID does not match **sem_perm.lclgid** 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 a `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 is true:

The process's effective user ID is super-user.

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

The process's effective user ID does not match `shm_perm.lcuid` and the process's effective group ID matches `shm_perm.lcgid` and the appropriate bit of the "group" portion (060) of `shm_perm.mode` is set.

The process's effective user ID does not match `shm_perm.lcuid` and the process's effective group ID does not match `shm_perm.lcgid` and the appropriate bit of the "other" portion (06) of `shm_perm.mode` is set.

Otherwise, the corresponding permissions are denied.

SEE ALSO
intro(3).

NAME

accept — accept a connection on a socket

SYNOPSIS

```
#include <net/socket.h>
```

```
accept(s, from)
```

```
int s;
```

```
struct sockaddr *from;
```

DESCRIPTION

This call is used to *accept* a connection on socket *s*; *from* is a result value indicating the address of the entity which connected, as known to the communications layer. This call is used with connection-based socket types, currently with SOCK_STREAM.

If the underlying communications layer has already made a connection on the socket, then the call returns immediately. If no connection has yet been made and the socket is nonblocking (see *ioctl(2)*), then a *-1* is returned and the global variable *errno* is set to EWOULDBLOCK. It is possible to *select(2)* a socket for the purposes of doing an *accept* by selecting it for read, since no data may be read until the connection completes.

SEE ALSO

connect(2), select(2), socket(2).

DIAGNOSTICS

Zero is returned if a connection is accepted; *-1* is returned in the error cases. Some important errors returned in *errno* are EOPNOTSUPP if the socket is not of a type supporting this operation, and EISCONN if the socket is already connected.

BUGS

This call is provisional and will exist in a slightly different form in future releases.

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:

A component of the path prefix is not a directory. [ENOTDIR]

Read, write, or execute (search) permission is requested for a null path name. [ENOENT]

The named file does not exist. [ENOENT]

Search permission is denied on a component of the path prefix. [EACCES]

Write access is requested for a file on a read-only file system. [EROFS]

Write access is requested for a pure procedure (shared text) file that is being executed. [ETXTBSY]

Permission bits of the file mode do not permit the requested access. [EACCES]

Path points outside the process's allocated address space. [EFAULT]

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.

Notice that it is only access bits that are checked. A directory may be announced as writable by *access*, but an attempt to open it for writing will fail because it is not allowed to write into the directory structure itself, although files may be created there. A file may look executable, but *exec* will fail unless it is in proper format.

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).

ASSEMBLER

```
moveq #33,D0
```

```

movl  path,A0
movl  amode,D1
trap  #0

```

Carry bit set on failure and cleared on success.

NAME

acct — enable or disable process accounting

SYNOPSIS

```

int acct (path)
char *path;

```

DESCRIPTION

Acct is used to enable or disable the system's process accounting routine. If the routine is enabled, an accounting record will be written on an accounting file for each process that terminates. Termination can be caused by one of two things: an *exit* call or a signal; see *exit(2)* and *signal(2)*. The effective user ID of the calling process must be super-user to use this call.

Path points to a path name naming the accounting file. The accounting file format is given in *acct(4)*.

The accounting routine is enabled if *path* is non-zero and no errors occur during the system call. It is disabled if *path* is zero and no errors occur during the system call.

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

The effective user ID of the calling process is not super-user. [EPERM]

An attempt is being made to enable accounting when it is already enabled. [EBUSY]

A component of the path prefix is not a directory. [ENOTDIR]

One or more components of the accounting file's path name do not exist. [ENOENT]

A component of the path prefix denies search permission. [EACCES]

The file named by *path* is not an ordinary file. [EACCES]

Mode permission is denied for the named accounting file. [EACCES]

The named file is a directory. [EISDIR]

The named file resides on a read-only file system. [EROFS]

Path points to an illegal address. [EFAULT]

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

acct(4).

ASSEMBLER

```

moveq #51,D0
movl  path,A0
trap  #0

```

Carry bit set on failure and cleared on success.

NAME

alarm — set a process's alarm clock

SYNOPSIS

```
unsigned alarm (sec)
unsigned sec;
```

DESCRIPTION

Alarm instructs the calling process's alarm clock 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 calling process's alarm clock. If the argument is 0, any alarm request is canceled. Because the clock has a 1-second resolution, the signal may occur up to one second early; because of scheduling delays, resumption of execution of when the signal is caught may be delayed an arbitrary amount. The longest specifiable delay time is 4,294,967,295 ($2^{32}-1$) seconds, or 136 years.

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

RETURN VALUE

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

SEE ALSO

pause(2), signal(2).

ASSEMBLER

```
moveq #27,D0
movl  sec,A0
trap  #0
```

On return, **D0** will contain the amount of time previously remaining in the alarm clock.

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; 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 location beyond the end of the data segment. 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.

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 break value being greater than or equal to the start address of any attached shared memory segment (see *shmop(2)*).

RETURN VALUE

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

SEE ALSO

exec(2).

ASSEMBLER

```
moveq #17,D0
movl  endds,A0
trap  #0
```

Carry bit cleared if the *brk* could be set; *brk* fails if the program requests more memory than the system limit or, on memory management CPUs, if too many segmentation registers would be required to implement the break.

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:

A component of the path name is not a directory. [ENOTDIR]

The named directory does not exist. [ENOENT]

Search permission is denied for any component of the path name. [EACCES]

Path points outside the process's allocated address space. [EFAULT]

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).

ASSEMBLER

```
moveq #12,D0
movl  path,A0
trap  #0
```

Carry bit set on failure and cleared on success.

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 (or 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 or 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 (see the *cc -n* option), 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.

Changing the owner of a file turns off the set-user-id bit, unless the superuser does it. This makes the system somewhat more secure at the expense of a degree of compatibility.

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

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

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).

ASSEMBLER

```
moveq #15,D0
movl  path,A0
movl  mode,D1
trap  #0
```

Carry bit set on failure and cleared on success.

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:

A component of the path prefix is not a directory. [ENOTDIR]

The named file does not exist. [ENOENT]

Search permission is denied on a component of the path prefix. [EACCES]

The effective user ID does not match the owner of the file and the effective user ID is not super-user. [EPERM]

The named file resides on a read-only file system. [EROFS]

Path points outside the process's allocated address space. [EFAULT]

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).

ASSEMBLER

```
moveq #16,D0
movl  path,A0
movl  owner,D1
movl  group,A1
trap  #0
```

Carry bit set on failure and cleared on success.

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 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, *..* can not 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:

Any component of the path name is not a directory. [ENOTDIR]

The named directory does not exist. [ENOENT]

The effective user ID is not super-user. [EPERM]

Path points outside the process's allocated address space. [EFAULT]

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).

ASSEMBLER

```
moveq #61,D0
movl  path,A0
trap  #0
```

Carry bit set on failure and cleared on success.

NAME

close — close a file descriptor

SYNOPSIS

```
int close (fildes)
int fildes;
```

DESCRIPTION

Fildes is a file descriptor obtained from a *creat*, *open*, *dup*, *fcntl*, or *pipe* system call. *Close* closes the file descriptor indicated by *fildes*. A close of all files is automatic on *exit*, but since there is a 20 open file limit on the number of open files per process, *close* is necessary for programs which deal with many files.

Close will fail if *fildes* is not a valid open file descriptor. [EBADF]

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).

ASSEMBLER

```
moveq #6,D0
movl fildes,A0
trap #0
```

Carry bit set on failure and cleared on success.

NAME

connect — initiate a connection on a socket

SYNOPSIS

```
#include <net/socket.h>
```

```
connect(s, addr)
int s;
struct sockaddr *addr;
```

DESCRIPTION

Connect causes a connection request to be initiated to the entity at *addr* using the underlying protocol of the socket *s*. When the connection completes, a zero value is returned.

If the socket is non-blocking but the connection cannot be completed immediately, then the call returns -1 and sets the external variable *errno* to EWOULDBLOCK. It is possible to *select*(2) a socket which is connecting by selecting it for writing, since writing is not possible before the connection completes.

If the socket is already connected, a value of -1 is returned and *errno* is set to EISCONN. Failure to connect often results in ETIMEDOUT or EREFUSED errors. Other errors are also possible.

SEE ALSO

accept(2), select(2), socket(2).

BUGS

A socket's state is not properly restored if a *connect* fails; for the time being you can *close* the socket and recreate it to get around the bug.

This call is provisional and will exist in a slightly different form in future releases.

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 process's effective user ID, the file's group ID is set to the process's effective group ID, 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, a non-negative integer, namely 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.

The *mode* given is arbitrary; it need not allow writing. This feature is used by programs which deal with temporary files of fixed names. The creation is done with a mode that forbids writing. Then, if a second instance of the program attempts a *creat*, an error is returned and the program knows that the name is unusable for the moment.

The system-scheduling algorithm does not make this a true uninterruptible operation, and a race condition may develop if *creat* is done at precisely the same time by two different processes.

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

A component of the path prefix is not a directory. [ENOTDIR]

A component of the path prefix does not exist. [ENOENT]

Search permission is denied on a component of the path prefix. [EACCES]

The path name is null. [ENOENT]

The file does not exist and the directory in which the file is to be created does not permit writing. [EACCES]

The named file resides or would reside on a read-only file system. [EROFS]

The file is a pure procedure (shared text) file that is being executed. [ETXTBSY]

The file exists and write permission is denied. [EACCES]

The named file is an existing directory. [EISDIR]

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

Path points outside the process's allocated address space. [EFAULT]

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

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

ASSEMBLER

```
moveq #8,D0
movl  path,A0
movl  mode,D1
trap  #0
```

Carry bit set on failure and cleared on success.

The file descriptor is returned in **D0**.

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:

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

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

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

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

SYNOPSIS

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

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

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

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

int execlp (file, arg0, arg1, ..., argn, 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)*), a text segment, and a data segment. The data segment contains 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.

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)*). The shell is invoked if a command file is found by *execlp* or *execvp*.

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 and is directly usable in another *execv* because *argv[argc]* is 0.

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 calling process's environment in the global cell:

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

and it is used to pass the calling process's environment 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.

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)*)

From C, two interfaces are available. *execl* is useful when a known file with known arguments is being called; the arguments to *execl* are the character strings constituting the file and the arguments; the first argument is conventionally the same as the file name (or its last component). A 0 argument must end the argument list.

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.

Envp is a pointer to an array of strings that constitute the *environment* of the process. Each string consists of a name, an =, and a null-terminated value. The array of pointers is terminated by a null pointer. The shell *sh(1)* passes an environment entry for each global shell variable defined when the program is called. See *environ(5)* for some conventionally used names. The C run-time start-off routine places a copy of *envp* in the global cell *environ*, which is used by *execv* and *execl* to pass the environment to any subprograms executed by the current program. The *exec* routines use lower-level routines as follows to pass an environment explicitly:

```
execve(file, argv, environ);
execl(file, arg0, arg1, . . . , argn, 0, environ);
```

Execlp and *execvp* are called with the same arguments as *execl* and *execv*, but duplicate the shell's actions in searching for an executable file in a list of directories. The directory list is obtained from the environment.

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

- One or more components of the new process file's path name do not exist. [ENOENT]

- A component of the new process file's path prefix is not a directory. [ENOTDIR]

- 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. [EACCES]

- 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. [ENOEXEC]

- The new process file is a pure procedure (shared text) file that is currently open for writing by some process. [ETXTBSY]

- The new process requires more memory than is allowed by the system-imposed maximum MAXMEM. [ENOMEM]

- The number of bytes in the new process's argument list is greater than the system-imposed limit of 5120 bytes. [E2BIG]

- 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. [EFAULT]

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

exit(2), *fork(2)*, *environ(5)*.

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 processes 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

signal(2), *wait(2)*.

WARNING

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

ASSEMBLER

```
moveq #1,D0
movl  status,A0
trap  #0
```

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 *cmds* available are:

- F_DUPFD** Return a new file descriptor as follows:
 Lowest numbered available file descriptor greater than or equal to *arg*.
 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 *exec(2)* system calls.
- F_GETFD** Get the close-on-exec flag associated with the file descriptor *fildes*. If the low-order bit is 0 the file will remain open across *exec*, otherwise the file will be closed upon execution of *exec*.
- F_SETFD** Set the close-on-exec flag associated with *fildes* to the low-order bit of *arg* (0 or 1 as above).
- F_GETFL** Get *file* status flags.
- F_SETFL** Set *file* status flags to *arg*. Only certain flags can be set; see *fcntl(5)*.

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

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

Cmd is F_DUPFD and 20 file descriptors are currently open. [EMFILE]

Cmd is F_DUPFD and *arg* is negative or greater than 20. [EINVAL]

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.

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

SEE ALSO

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

ASSEMBLER

```

moveq #62,D0
movl  fildes,A0
movl  cmd,D1
movl  arg,A1
trap  #0

```

Carry bit set on failure and cleared on success.

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 (see *times(2)*).

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

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. [EAGAIN]

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.

FORK (2)

SEE ALSO

exec(2), times(2), wait(2).

ASSEMBLER

```
moveq #2,D0
trap #0
```

New process return.

Old process return, new process ID in D0.

Carry bit cleared on success.

The return locations in the old and new process differ by one 16 bit word.

The C-bit is set in the old process if a new process could not be created.

FORK (2)

GETHOSTNAME(2)

(UniSoft)

GETHOSTNAME(2)

NAME

gethostname — get name of current host

SYNOPSIS

```
char hostname[32];
```

```
gethostname(hostname, sizeof (hostname));
```

DESCRIPTION

Gethostname returns the standard host name for the current processor, as set by *sethostname(2)* and defined in *rhost(3)*. The name is null-terminated.

SEE ALSO

sethostname(2), rhost(3).

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.

These system calls are useful for generating uniquely-named temporary files.

SEE ALSO

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

ASSEMBLER

```
moveq #20,D0      | getpid
trap   #0
```

Process ID is returned in D0.

```
moveq #39,D0      | getpgrp
movl  #0,A0
trap   #0
```

Process ID is returned in D0.

```
moveq #20,D0      | getppid
trap   #0
```

Parent process ID is returned in D1.

NAME

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

SYNOPSIS

```
int getuid ()
int geteuid ()
int getgid ()
int 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).

ASSEMBLER

```
moveq #24,D0      | sys getuid
trap   #0
```

Real user ID returned in D0.

```
moveq #24,D0      | sys geteuid
trap   #0
```

Effective user ID returned in D1.

```
moveq #47,D0      | sys getgid
trap   #0
```

Real group ID returned in D0.

```
moveq #47,D0      | sys getegid
trap   #0
```

Effective group ID returned in D1.

NAME

`ioctl` — control device

SYNOPSIS

`ioctl (fildes, request, arg)`

DESCRIPTION

`ioctl` performs a variety of functions on character special files (devices). The writeups of various devices in Section 7 discuss how `ioctl` applies to them.

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

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

Fildes is not associated with a character special device. [ENOTTY]

Request or *arg* is not valid. See Section 7. [EINVAL]

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 *UniPlus+ Administrator's Manual*.

ASSEMBLER

```

moveq #54,D0      | sys ioctl
movl  fildes,A0
movl  request,D1
movl  #argp,A1
trap  #0

```

Carry bit cleared on success.

```

moveq #31,D0      | sys stty
movl  fildes,A0
movl  #argp,D1
trap  #0

```

Carry bit cleared on success.

```

moveq #32,D0      | sys gtty
movl  fildes,A0
movl  request,D1
movl  arg,A1
trap  #0

```

Carry bit set on failure and cleared on success.

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, or the process is sending to itself.

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:

Sig is not a valid signal number. [EINVAL]

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

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

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

kill(1), getpid(2), setpgid(2), signal(2).

ASSEMBLER

```

moveq #37,D0
movl  pid,A0
movl  sig,D1
trap  #0

```

Carry bit set on failure and cleared on success.

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:

- A component of either path prefix is not a directory. [ENOTDIR]
- A component of either path prefix does not exist. [ENOENT]
- A component of either path prefix denies search permission. [EACCES]
- The file named by *path1* does not exist. [ENOENT]
- The link named by *path2* exists. [EEXIST]
- The file named by *path1* is a directory and the effective user ID is not super-user. [EPERM]
- The link named by *path2* and the file named by *path1* are on different logical devices (file systems). [EXDEV]
- Path2* points to a null path name. [ENOENT]
- The requested link requires writing in a directory with a mode that denies write permission. [EACCES]
- The requested link requires writing in a directory on a read-only file system. [EROFS]
- Path* points outside the process's allocated address space. [EFAULT]
- The requested link requires the file named by *path1* to have more than 1000 links. [EMLINK]

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).

ASSEMBLER

```

moveq #9,D0
movl  path1,A0
movl  path2,D1
trap  #0

```

Carry bit set on failure and cleared on success.

NAME

lockf — provide exclusive file regions for reading or writing

SYNOPSIS

lockf(fildes, mode, size)

int fildes;

int mode;

int size;

DESCRIPTION

Lockf will allow a specified number of bytes to be accessed only by the locking process. Other processes which attempt to lock, read, or write the locked area will sleep until the area becomes unlocked.

Fildes is the word returned from a successful *open*, *creat*, *dup*, or *pipe* system call.

Mode is zero to unlock the area. *Mode* is one or two for making the area locked. If the *mode* is one and the area has some other lock on it, then the process will sleep until the entire area is available. If the *mode* is two and the area is locked, an error will be returned.

Size is the number of contiguous bytes to be locked or unlocked. The area to be locked starts at the current offset in the file. If *size* is zero, the area to the end of file is locked.

The potential for a deadlock occurs when a process controlling a locked area is put to sleep by accessing another process's locked area. Thus calls to *lockf*, *read*, or *write* scan for a deadlock prior to sleeping on a locked area. An error return is made if sleeping on the locked area would cause a deadlock.

Lock requests may, in whole or part, contain or be contained by a previously locked area for the same process. When this or adjacent areas occur, the areas are combined into a single area. If the request requires a new lock element with the lock table full, an error is returned, and the area is not locked.

Unlock requests may, in whole or part, release one or more locked regions controlled by the process. When regions are not fully released, the remaining areas are still locked by the process. Release of the center section of a locked area requires an additional lock element to hold the cut off section. If the lock table is full, an error is returned, and the requested area is not released.

While locks may be applied to special files or pipes, read/write operations will not be blocked. Locks may not be applied to a directory.

SEE ALSO

close(2), creat(2), dup(2), open(2), read(2), write(2).

DIAGNOSTICS

The value -1 is returned if the file does not exist, or if a deadlock using file locks would occur. EACCES will be returned for lock requests in which the area is already locked by another process. EDEADLOCK will be returned by: read, write, or locking if a deadlock would occur. EDEADLOCK will also be returned when the locktable overflows.

ASSEMBLER

```

moveq #56,D0
movl  fildes,A0
movl  mode,D1
movl  size,A1
trap  #0

```

Carry bit cleared on success.

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:

Fildes is not an open file descriptor. [EBADF]

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

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

The resulting file pointer would be negative. [EINVAL]

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)`.

ASSEMBLER

```

moveq #19,D0
movl  fildes,A0
movl  offset,D1
movl  whence,A1
trap  #0

```

Carry bit set on failure and cleared on success.

File offset returned in `D0`.

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

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:

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

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

mkdir(1), chmod(2), exec(2), umask(2), fs(4).

ASSEMBLER

```

moveq #14,D0
movl  path,A0
movl  mode,D1
movl  dev,A1
trap  #0

```

Carry bit set on failure and cleared on success.

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. Physically write-protected and magnetic tape file systems must be mounted read-only or errors will occur when access times are updated, whether or not any explicit write is attempted.

Mount may be invoked only by the super-user.

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

The effective user ID is not super-user. [EPERM]

Any of the named files does not exist. [ENOENT]

A component of a path prefix is not a directory. [ENOTDIR]

Spec is not a block special device. [ENOTBLK]

The device associated with *spec* does not exist. [ENXIO]

Dir is not a directory. [ENOTDIR]

Spec or *dir* points outside the process's allocated address space. [EFAULT]

Dir is currently mounted on, is someone's current working directory or is otherwise busy. [EBUSY]

The device associated with *spec* is currently mounted. [EBUSY]

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).

ASSEMBLER

```

moveq #21,D0      | sys mount
movl  spec,A0
movl  dir,D1
movl  rwflag,A1
trap  #0

```

Carry bit set on failure and cleared on success.

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:

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

Cmd is not a valid command. [EINVAL]

Cmd is equal to `IPC_STAT` and {READ} operation permission is denied to the calling process (see *intro(2)*). [EACCESS]

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 `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. [EPERM]

Buf points to an illegal address. [EFAULT]

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

msgget(2), msgop(2).

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:

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_lrpil`, `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:

A message queue identifier exists for *key* but operation permission (see *intro(2)*) as specified by the low-order 9 bits of *msgflg* would not be granted. [EACCES]

A message queue identifier does not exist for *key* and $(msgflg \& IPC_CREAT)$ is "false". [ENOENT]

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. [ENOSPC]

A message queue identifier exists for *key* but $((msgflg \& IPC_CREAT) \& (msgflg \& IPC_EXCL))$ is "true". [EEXIST]

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

msgctl(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:

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

Operation permission is denied to the calling process (see *intro(2)*). [EACCESS]

Mtype is less than 1. [EINVAL]

The message cannot be sent for one of the reasons cited above and (*msgflg* & IPC_NOWAIT) is "true". [EAGAIN]

Msgsz is less than zero or greater than the system imposed limit. [EINVAL]

Msgp points to an illegal address. [EFAULT]

Upon successful completion, the following actions are taken with respect to the data structure associated with *msqid* (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 *msqid* 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.

Msqid 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:

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

Operation permission is denied to the calling process. [EACCESS]

Msgsz is less than 0. [EINVAL]

Mtext is greater than *msgsz* and (*msgflg* & MSG_NOERROR) is "false". [E2BIG]

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

Msgp points to an illegal address. [EFAULT]

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

msgctl(2), *msgget(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.

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

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

nice(1), exec(2).

ASSEMBLER

```
moveq #34,D0
movl  incr,A0
trap  #0
```

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_CREAT If the file exists, this flag has no effect. Otherwise, the file's owner ID is set to the process's effective user ID, the file's group ID is set to the process's effective group ID, 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 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)*.

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.

Upon successful completion a non-negative integer, the file descriptor, is returned.

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)*.

No process may have more than 20 file descriptors open simultaneously.

The named file is opened unless one or more of the following are true:

A component of the path prefix is not a directory. [ENOTDIR]

O_CREAT is not set and the named file does not exist. [ENOENT]

A component of the path prefix denies search permission. [EACCESS]

Oflag permission is denied for the named file. [EACCESS]

The named file is a directory and *oflag* is write or read/write. [EISDIR]

The named file resides on a read-only file system and *oflag* is write or read/write. [EROFS]

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

The named file is a character special or block special file, and the device associated with this special file does not exist. [ENXIO]

The file is a pure procedure (shared text) file that is being executed and *oflag* is write or read/write. [ETXTBSY]

Path points outside the process's allocated address space. [EFAULT]

O_CREAT and **O_EXCL** are set, and the named file exists. [EEXIST]

O_NDELAY is set, the named file is a FIFO, **O_WRONLY** is set, and no process has the file open for reading. [ENXIO]

RETURN VALUE

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

SEE ALSO

close(2), *creat(2)*, *dup(2)*, *fcntl(2)*, *lseek(2)*, *read(2)*, *write(2)*.

ASSEMBLER

```
moveq #5,D0
movl  path,A0
movl  oflag,D1
movl  mode,A1
trap  #0
```

Carry bit set on failure and cleared on success.

File descriptor is returned in **D0**.

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)*.

ASSEMBLER

```
moveq #29,D0
trap  #0
```

NAME

phys — allow a process to access physical addresses

SYNOPSIS

```
phys(physnum, virtaddr, size, physaddr)
int physnum
char *virtaddr;
long size;
char *physaddr;
```

DESCRIPTION

The *phys(2)* call maps arbitrary physical memory into a process's virtual address space. *Physnum* is a number (0-3) that specifies which of 4 physical spaces to set up. Up to 4 *phys(2)* calls can be active at any one time. *Virtaddr* is the process's virtual address. *Size* is the number of bytes to map in. *Physaddr* is the physical address to map in.

Valid *virtaddr* and *physaddr* values are constrained by hardware and must be at an address multiple of the resolution of the CPU's memory management scheme. If *size* is non zero, size is rounded up to the next MMU resolution boundary. If *size* is zero, any previous *phys(2)* mapping for that *physnum* segment is nullified.

For example, the call:

```
phys(2, 0x100000, 32768, 0)
```

will allow a process to access physical locations 0 through 32767 by referencing virtual address 0x100000 through 0x100000+32767.

In actuality, the CPU MMU register is loaded with *physaddr* shifted to account for page resolution.

Phys(2) may only be executed by the super-user.

DIAGNOSTICS

The value zero is returned if the *phys* call was successful. The value -1 is returned if not super-user, if *virtaddr* or *physaddr* is not in the proper range, or if the specified *virtaddr* segment register is already in use.

BUGS

This system call is very machine dependent.

ASSEMBLER

```
moveq #55,D0
movl  physnum,A0
movl  virtaddr,D1
movl  size,A1
movl  D2,save
movl  physaddr,D2
trap  #0
movl  save,D2
```

Carry bit cleared on success.

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.

Writes up to 5120 bytes of data are buffered by the pipe before the writing process is blocked. A read on file descriptor *fildes*[0] accesses the data written to *fildes*[1] on a first-in-first-out basis.

No process may have more than 20 file descriptors open simultaneously.

Pipe will fail if 19 or more file descriptors are currently open. [EMFILE]

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

sh(1), read(2), write(2).

ASSEMBLER

```
moveq #42,D0
movl  fildes,A0
trap  #0
```

Carry bit set on failure and cleared on success.

Read file descriptor in **D0**. Write file descriptor in **D1**.

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 & data segments into memory (process lock)

TXTLOCK —

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:

The effective user ID of the calling process is not super-user. [EPERM]

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 **TXTLOCK** 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. [EINVAL]

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)*.

ASSEMBLER

```
movq  #45,D0
```

```
movl  op,A0
```

```
trap  #0
```

NAME

profil — execution time profile

SYNOPSIS

```

profil (buff, bufsiz, offset, scale)
char *buff;
int bufsiz, offset, scale;

```

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; *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 (16 bit), fixed-point fraction with binary point at the left: FFFF (hex) gives a 1-1 mapping of *pc*'s to words in *buff*; FFFF (hex) maps each pair of instruction words together. 2(hex) 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

prof(1), monitor(3C).

ASSEMBLER

```

moveq #44,D0
movl  buff,A0
movl  bufsiz,D1
movl  offset,A1
movl  D2,save
movl  scale,D2
trap  #0
movl  save,D2

```

The D2 register must be saved when calling *profil*(2) since that register might be in use by the "C" program that calls this routine.

NAME

ptrace — process trace

SYNOPSIS

```
int ptrace (request, pid, addr, data);
int request, pid, addr, data;
```

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. 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. 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 are system dependent. 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*. 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
 - the condition codes
 - the floating point status register and floating point registers
 - certain bits of the Processor Status Word
- 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 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.
- Note: the trace bit remains set after an interrupt.

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:

Request is an illegal number. [EIO]

Pid identifies a child that does not exist or has not executed a *ptrace* with request 0. [ESRCH]

SEE ALSO

exec(2), *signal(2)*, *wait(2)*.

ASSEMBLER

```

moveq #26,D0
movl  D2,save | save D2 register
crl  _errno
movl  request,A0
movl  pid,D1
movl  addr,A1
movl  data,D2
trap  #0
movl  save,D2 | restore D2 register

```

Carry bit set on failure and cleared on success.

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:

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

Buf points outside the allocated address space. [EFAULT]

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)*.

ASSEMBLER

```

moveq #3,D0
movl  fildes,A0
movl  buf,D1
movl  nbytes,A1
trap  #0

```

NAME

reboot - reboot the system

SYNOPSIS

reboot ()

DESCRIPTION

Reboot causes the kernel to execute the initial bootstrap code that was used to boot the operating system.

On most CPUs the *reboot(2)* command will take the place of a manual restart.

ASSEMBLER

```
moveq 64,D0
trap  #0
```

NAME

receive — receive message from a socket

SYNOPSIS

```
#include <net/socket.h>
```

```
cc = receive(s, from, buf, len);  
int cc, s;  
struct sockaddr *from;  
char *buf;  
int len;
```

DESCRIPTION

Receive is used to receive a message from a SOCK_DGRAM or SOCK_RAW socket. The source address of the message is placed in *from*. The length of the message is returned in *cc*. If the message is too long to fit in the supplied buffer, then excess characters are discarded.

If no messages are available at the socket, the *receive* waits for a message to arrive, unless the socket is nonblocking in which case a *cc* of *-1* is returned with the external variable *errno* set to *EWouldBlock*.

The *select(2)* call may be used to determine when more data arrives.

SEE ALSO

send(2), *socket(2)*.

BUGS

This call is provisional and will exist in a slightly different form in future releases.

NAME

select — synchronous i/o multiplexing

SYNOPSIS

```
nfd = select(nfds, readfds, writefds, milli);  
int nfds;  
int *readfds, *writefds;  
int milli;
```

DESCRIPTION

Select examines the i/o descriptors specified by the bit masks *readfds* and *writefds* to see if they are ready for reading and/or writing respectively and returns, in place, a mask of those descriptors which are ready. The total number of ready descriptors is returned in *nfd*.

Milli is the maximum number of milliseconds to wait before giving up if no descriptors come active. If no maximum wait is desired a very large integer can be given.

A *milli* of 0 specifies a poll; the *select* returns whatever information is available without blocking. Either *readfds* or *writefds* may be given as 0 if no descriptors are interesting.

For the present, since UNIX allows only 20 file descriptors it suffices for *nfd* to be 20, and for *readfds* and *writefds* to be pointers to integer variables. File descriptor *f* is represented by the bit "1 < f" in the mask.

SEE ALSO

accept(2), connect(2), ioctl(2), read(2), receive(2), send(2), write(2).

BUGS

The system currently rounds *milli* to integral seconds, with a resolution of ± 1 second.

Currently *select* only works correctly on sockets and psuedo-teletypes. Other file-descriptors always *select* as ready.

This call is provisional and will exist in a slightly different form in future releases.

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 *semval* (see *intro(2)*). {READ}

SETVAL Set the value of *semval* to *arg.val*. {ALTER} When this *cmd* is successfully executed the *semadj* value corresponding to the specified semaphore in all processes is cleared.

GETPID Return the value of *sempid*. {READ}

GETNCNT Return the value of *semncnt*. {READ}

GETZCNT Return the value of *semzcnt*. {READ}

The following *cmds* return and set, respectively, every *semval* in the set of semaphores.

GETALL Place *semvals* into array pointed to by *arg.array*. {READ}

SETALL Set *semvals* according to the array pointed to by *arg.array*. {ALTER} When this *cmd* is successfully executed the *semadj* 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 *semid* into the structure pointed to by *arg.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 *semid* to the corresponding value found in the structure pointed to by *arg.buf*:

```
sem_perm.uid
sem_perm.gid
sem_perm.mode /* 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:

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. [EINVAL]

Operation permission is denied to the calling process (see *intro(2)*). [EACCESS]

Cmd is SETVAL or SETALL and the value to which `semval` is to be set is greater than the system imposed maximum. [ERANGE]

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`. [EPERM]

Arg.buf points to an illegal address. [EFAULT]

RETURN VALUE

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

GETVAL	The value of <code>semval</code> .
GETPID	The value of <code>sempid</code> .
GETNCNT	The value of <code>semncnt</code> .
GETZCNT	The value of <code>semzcnt</code> .
All others	A value of 0.

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

SEE ALSO

`semget(2)`, `semop(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:

Nsems is either less than or equal to zero or greater than the system imposed limit. [EINVAL]

A semaphore identifier exists for *key* but operation permission (see *intro(2)*) as specified by the low-order 9 bits of `semflg` would not be granted. [EACCESS]

A semaphore identifier exists for *key* but the number of semaphores in the set associated with it is less than *nsems* and *nsems* is not equal to zero. [EINVAL]

A semaphore identifier does not exist for *key* and (`semflg & IPC_CREAT`) is "false". [ENOENT]

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. [ENOSPC]

A semaphore identifier exists for *key* but (`(semflg & IPC_CREAT) & (semflg & IPC_EXCL)`) is "true". [EEXIST]

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

`semctl(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 atomically 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 occurs:

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*:

Semid is not a valid semaphore identifier. [EINVAL]

Sem_num is less than zero or greater than or equal to the number of semaphores in the set associated with *semid*. [EFBIG]

Nsops is greater than the system imposed maximum. [E2BIG]

Operation permission is denied to the calling process (see *intro(2)*). [EACCES]

The operation would result in suspension of the calling process but (*sem_flg* & IPC_NOWAIT) is "true". [EAGAIN]

The limit on the number of individual processes requesting an SEM_UNDO would be exceeded. [ENOSPC]

The number of individual semaphores for which the calling process requests a SEM_UNDO would exceed the limit. [EINVAL]

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. [ERANGE]

Sops points to an illegal address. [EFAULT]

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.

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)*, *semctl(2)*, *semget(2)*.

NAME

send — send message from a socket

SYNOPSIS

```
#include <net/socket.h>
```

```
send(s, to, msg, len)
int cc, s;
struct sockaddr *to;
char *msg;
int len;
```

DESCRIPTION

Send is used to transmit a message to another socket from a SOCK_DGRAM or SOCK_RAW socket. The address of the target is given by *to*. The length of the message is given by *len*. If the message is too long to pass atomically through the underlying protocol, then the error EMSGSIZE is returned, and the message is not transmitted.

No indication of failure to deliver is implicit in *send*. Some locally detected errors may be reported to the user through the return value from *send* being -1 with the errors being stored in the external variable *errno*.

If no messages space is available at the socket to hold the message to be transmitted, then *send* normally blocks, unless the socket is SO_NONBLOCKING in which case a *cc* of -1 is returned with the external variable *errno* set to EWOULDBLOCK. The *select(2)* call may be used to determine when it is possible to send more data.

SEE ALSO

send(2), socket(2).

BUGS

This call is provisional and will exist in a slightly different form in future releases.

NAME

sethostname — set name of host cpu

SYNOPSIS

```
sethostname(name, namelen)
char *name;
int namelen;
```

DESCRIPTION

This call sets the name of the host processor to be *name*, which has length *namelen* characters. This is normally executed when the system is bootstrapped, executed out of the file */etc/rc*. The name set should not be a nickname for the machine, but the full name of the machine, i.e., "unisoft".

SEE ALSO

gethostname(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).

ASSEMBLER

```
moveq #39,D0  
movw #1,A0  
trap #0
```

Carry bit set on failure and cleared on success.

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)*.

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]

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).

ASSEMBLER

```
moveq #23,D0      | sys setuid
movl  uid,A0
trap  #0
```

Carry bit cleared on success.

```
moveq #46,D0      | sys setgid
movl  gid,A0
trap  #0
```

Carry bit set on failure and cleared on success.

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 shmids *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 *shmid* 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 *shmid* to the corresponding value found in the structure pointed to by *buf*:

```
shm_perm.uid
shm_perm.gid
shm_perm.mode /* 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 **shm_perm.uid** in the data structure associated with *shmid*.

- IPC_RMID** Remove the shared memory identifier specified by *shmid* from the system and destroy the shared memory segment 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 **shm_perm.uid** in the data structure associated with *shmid*.

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

Shmid is not a valid shared memory identifier. [EINVAL]

Cmd is not a valid command. [EINVAL]

Cmd is equal to **IPC_STAT** and [READ] operation permission is denied to the calling process (see *intro(2)*). [EACCESS]

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 **shm_perm.uid** in the data structure associated with *shmid*. [EPERM]

Buf points to an illegal address. [EFAULT]

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

shmget(2), shmop(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:

Size is less than the system-imposed minimum or greater than the system-imposed maximum. [EINVAL]

A shared memory identifier exists for *key* but operation permission (see *intro(2)*) as specified by the low-order 9 bits of *shmflg* would not be granted. [EACCESS]

A shared memory identifier exists for *key* but the size of the segment associated with it is less than *size* and *size* is not equal to zero. [EINVAL]

A shared memory identifier does not exist for *key* and $(shmflg \& IPC_CREAT)$ is “false”. [ENOENT]

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. [ENOSPC]

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. [ENOMEM]

A shared memory identifier exists for *key* but $(shmflg \& IPC_CREAT) \& (shmflg \& IPC_EXCL)$ is “true”. [EEXIST]

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

`shmctl(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 $(shmaddr - (shmaddr \text{ modulus } SHMLBA))$.

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:

Shmid is not a valid shared memory identifier. [EINVAL]

Operation permission is denied to the calling process (see *intro(2)*). [EACCESS]

The available data space is not large enough to accommodate the shared memory segment. [ENOMEM]

Shmaddr is not equal to zero, and the value of $(shmaddr - (shmaddr \text{ modulus } SHMLBA))$ 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. [EINVAL]

The number of shared memory segments attached to the calling process would exceed the system-imposed limit. [EMFILE]

Shmdt detaches from the calling process's data segment the shared memory segment located at the address specified by *shmaddr*.

Shmdt will fail and not detach the shared memory segment if *shmaddr* is not the data segment start address of a shared memory segment. [EINVAL]

RETURN VALUES

Upon successful completion, the return value is as follows:

Shmat returns the data segment start address of the attached shared memory segment.

Shmctl 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)*, *shmctl(2)*, *shmget(2)*.

NAME

signal — specify what to do upon receipt of a signal

SYNOPSIS

```
#include <sys/signal.h>
int (*signal (sig, func))()
int sig;
int (*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 of 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 the following consequences:

All of the receiving process's open file descriptors will be closed.

If the parent process of the receiving process is executing a *wait*, it will be notified of the termination of the receiving process and the terminating signal's number will be made available to the parent process; see *wait(2)*.

If the parent process of the receiving process is not executing a *wait*, the receiving process will be transformed into a zombie process (see *exit(2)* for definition of zombie process).

The parent process ID of each of the receiving process's existing child processes and zombie processes will be 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 receiving 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 will be written on the accounting file if the system's accounting routine is enabled; see `acct(2)`.

If the receiving process's process ID, tty group ID, and process group ID are equal, the signal `SIGHUP` will be sent to all of the processes that have a process group ID equal to the process group ID of the receiving process.

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. 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 will 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 one or more of the following are true:

`Sig` is an illegal signal number, including `SIGKILL`. [EINVAL]

`Func` points to an illegal address. [EFAULT]

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(1)`, `kill(2)`, `pause(2)`, `ptrace(2)`, `wait(2)`, `setjmp(3C)`.

WARNING

Two other signals that behave differently than the signals described above exist in this release of the system; they are:

<code>SIGCLD</code>	18	death of a child (reset when caught)
<code>SIGPWR</code>	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 of 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)`.

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 proceeding 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.

BUGS

If a repeated signal arrives before the last one can be reset, there is no chance to catch it.

The type specification of the routine and its *func* argument are problematical.

The symbols *sighnd* and *sigtrap* are globally defined symbols used by *signal(2)* and are reserved words.

NAME

socket — create an endpoint for communication

SYNOPSIS

```
#include <net/socket.h>
```

```
s = socket(type, pf, addr, options);
int type;
struct sockproto *pf;
struct sockaddr *addr;
int options;
```

DESCRIPTION

Socket creates a communication endpoint and returns a descriptor, much like a file descriptor. The socket has the specified *type* which defines the semantics of communication. Currently defined types are SOCK_STREAM, for sequenced, reliable, two-way connection based streams with an out-of-band mechanism; SOCK_DGRAM for datagrams, connectionless, unreliable messages of a fixed (typically small) maximum length, SOCK_RAW providing access to internal network interfaces. The type SOCK_RAW, which is available only to the super-user, is not described here.

The *pf* supplied causes a specific protocol to be used with the socket; since there is currently only one protocol supporting each socket type we will not discuss this further.

The *addr* parameter specifies the address for the socket. A socket address is a discriminated union with a fixed length of 16 bytes. The first two bytes indicates the format of the remaining bytes. The only currently relevant variant is a *sockaddr_in*, an internet address. The first three fields of a variable of this type are AF_INET (indicating that the address is of the Address Family Internet, this is defined in *<net/socket.h>*), a 16 bit socket number to be used (see *<net/in.h>* for lists of well-known sockets), and a 32 bit host address. The socket number and host address are in network byte order.

If no address is specified, then the system will assign one at its convenience; currently it does this at connection time to simplify the routing decisions required of the connected socket. If the socket number is omitted, a unique socket number will be supplied. The socket numbers in the range 0 to IPPORT_RESERVED-1 are reserved for the super-user.

The procedure *rhost(3)* may be used to determine Internet host numbers, while *raddr(3)* converts addresses to standard host names.

Sockets of type SOCK_STREAM are full-duplex byte streams, similar to two-way pipes. A typical use of such a stream involves creation with *socket* and connection to another socket with a *connect(2)* call, followed by a sequence of *read* and *write* calls to exchange data, followed finally by a *close(2)*. Out-of-band data may also be transmitted as described below.

The protocol used to implement a SOCK_STREAM insures that data is not lost or duplicated. If a piece of data for which the peer protocol has buffer space cannot be successfully transmitted within a reasonable length of time (typically about 1 minute), then the connection is considered broken and calls will indicate error with -1 returns with ETIMEDOUT as the specific code in the global variable *errno*. The protocols optionally keep sockets "warm" by forcing transmissions roughly every minute in the absence of

other activity. An error is then indicated if no response can be elicited on an otherwise idle connection for a extended period (e.g., 5 minutes). A SIGPIPE signal is raised if a process *writes* on a broken stream; this causes naive processes, which do not handle the signal, to exit.

SOCK_DGRAM sockets allow sending of datagrams to correspondents named in *send(2)* calls. It is also possible to receive datagrams at such a socket with *receive(2)*

The primitive *socketaddr(2)* can be used to determine the address of a socket.

The options available on sockets are ored together in *options*, and are:

SO_DEBUG

Enable protocol tracing for this socket, to be used in protocol debugging.

SO_ACCEPTCONN

which must be used with SOCK_STREAM sockets which are to accept connections. Only sockets which indicate SO_ACCEPTCONN as a creation parameter may do *accept(2)*, and such sockets may not do *connect(2)*.

SO_DONTLINGER

which allows *close(2)* operations on a socket to complete immediately. Otherwise the system will block a process waiting for data to drain (or return EWOULDBLOCK if the socket is marked NONBLOCKING) when a close is attempted. See also the SIOCSLINGER *ioctl* below.

SO_KEEPAIVE

which causes keep alive to be used so as to time out dead connections. If this option is not specified, then timing out dead connections is the responsibility of the user process.

General *ioctls* which apply to sockets are:

SIOCDONE

indicating that the user is done receiving (if the integer parameter is 0), sending (if the integer parameter is 1) or both (if the parameter is 2) on the indicated socket. This is normally used to indicate an end-of-file on a SOCK_STREAM while continuing to read input.

SIOCSLINGER

sets the linger time to the number of seconds specified by the integer parameter. This is currently only partly implemented: linger time is either 0 or infinite (if non-zero).

SIOCGLINGER

returns the current linger time.

FIONBIO

takes an integer parameter saying whether non-blocking i/o is desired on the specified socket. Applies to sockets and specifies that operations are to return EWOULDBLOCK rather than blocking. A *select(2)* operation may be used to determine when i/o is possible without busy polling.

The out-of-band data facilities of the stream protocols are currently primitive, allowing the user to send a single byte of out-of-band data to the correspondent process. An SIOCSENDOOB *ioctl* takes as parameter the

address of the character to be sent as a parameter. This causes a SIGURG signal, indicating an urgent condition, to be raised in the correspondent process, and places a mark in the data stream after the last byte written before the out-of-band data was sent.

The SIOCSPGRP *ioctl* can be used to specify a process group to receive the SIGURG signal when the out-of-band data arrives. If the integer argument to SIOCSPGRP is negative, then it is taken to mean a single process rather than a process group, given by the absolute value of the argument. The SIOCGPGRP *ioctl* returns the current value of a sockets process group.

When a process receives a SIGURG signal it can enquire of each of its channels to see which ones have out-of-band data, by doing SIOCRCVOOB on each channel. This will return EINVAL if there is no out-of-band data currently available on that channel. If a channel has out-of-band data, a course of action might be to read in the input stream to the mark, which can be detected by SIOCATMARK which returns a 0 or a 1 into its integer parameter telling whether the read pointer is now at the mark. The system never returns bytes on both sides of a mark with a single read.

Facilities to provide the user with interrupts whenever i/o is possible on a specifiable set of channels are planned. This will allow interrupt-driven i/o processing similar to the out-of-band facilities.

SEE ALSO

accept(2), *connect(2)*, *receive(2)*, *select(2)*, *send(2)*, *socketaddr(2)*.

BUGS

This call is provisional and will exist in a slightly different form in future releases.

NAME

socketaddr — return address associated with a socket

SYNOPSIS

```
#include <set/socket.h>
```

```
socketaddr(s, addr)  
int s;  
struct sockaddr *addr;
```

DESCRIPTION

The address associated with the socket *s* is returned in *addr*. If *s* is not a socket, -1 is returned and an appropriate *errno* is returned.

SEE ALSO

socket (2).

BUGS

This call is provisional and will exist in a slightly different form in future releases.

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 system 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(2)

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

A component of the path prefix is not a directory. [ENOTDIR]

The named file does not exist. [ENOENT]

Search permission is denied for a component of the path prefix. [EACCES]

Buf or *path* points to an invalid address. [EFAULT]

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

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

Buf points to an invalid address. [EFAULT]

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), time(2), unlink(2).

ASSEMBLER

```
moveq #18,D0      | sys stat
movl  path,A0
movl  buf,D1
trap  #0
```

Carry bit set on failure and cleared on success.

```
moveq #28,D0      | sys fstat
movl  fildes,A0
movl  buf,D1
trap  #0
```

Carry bit set on failure and cleared on success.

STAT(2)**STIME(2)****STIME(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.

Stime will fail if the effective user ID of the calling process is not super-user. [EPERM]

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).

ASSEMBLER

```
moveq #25,D0
movl  tp,A0
trap  #0
```

Carry bit set on failure and cleared on success.

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*.

ASSEMBLER

```
moveq 36,D0  
trap #0
```

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.

Time will fail if *tloc* points to an illegal address. [EFAULT]

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).

ASSEMBLER

moveq #13,D0

trap #0

tstl tloc

beq 1\$

movl tloc,A0

@ movl D0,A0@

	time(0)?
	yes, return

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 is 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.

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.

Times will fail if *buffer* points to an illegal address. [EFAULT]

RETURN VALUE

Upon successful completion, *times* returns the elapsed real time, in 60ths 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)*.

ASSEMBLER

```
moveq  #43,D0
movl   buffer,A0
trap   #0
```

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 process's file size limit. 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 process's file size limit 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)*.

ASSEMBLER

```
moveq #63,D0
movl  cmd,A0
movl  newlimit,D1
trap  #0
```

Carry bit set on failure and cleared on success.

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.

The file mode creation mask is used whenever a file is created by *creat(2)*, *mknod(2)* or *open(2)*. The actual mode (see *chmod(2)*) of the newly-created file is the difference between the given mode and *cmask*. In other words, *cmask* shows the bits to be turned off when a new file is created.

The previous value of *cmask* is returned by the call. The value is initially 022, which is an octal 'mask' number representing the complement of the desired mode. '022' here means that no permissions are withheld from the owner, but write permission is forbidden to group and to others. Its complement, the mode of the file, would be 755. The file mode creation mask is inherited by child processes.

RETURN VALUE

The previous value of the file mode creation mask is returned.

SEE ALSO

mkdir(1), sh(1), chmod(2), creat(2), mknod(2), open(2).

ASSEMBLER

```
moveq #60,D0
movl  cmask,A0
trap  #0
```

The previous value of umask is returned to **D0**.

NAME

umount — unmount a file system

SYNOPSIS

```
int umount (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:

The process's effective user ID is not super-user. [EPERM]

Spec does not exist. [ENXIO]

Spec is not a block special device. [ENOTBLK]

Spec is not mounted. [EINVAL]

A file on *spec* is busy. [EBUSY]

Spec points outside the process's allocated address space. [EFAULT]

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).

ASSEMBLER

```
moveq #22,D0      | sys umount
movl  spec,A0
trap  #0
```

Carry bit set on failure and cleared on success.

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>:

```
struct utsname {
    char    sysname[9];
    char    nodename[9];
    char    release[9];
    char    version[9];
    char    machine[9];
};
extern struct utsname utsname;
```

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.

Uname will fail if *name* points to an invalid address. [EFAULT]

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).

ASSEMBLER

```
moveq    #57,D0
movl     name,A0      | fetch argument
subl     A1,A1        |  uname
trap     #0
```

Carry bit set on failure and cleared on success.

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:

A component of the path prefix is not a directory. [ENOTDIR]

The named file does not exist. [ENOENT]

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. [EACCES]

The named file is a directory and the effective user ID of the process is not super-user. [EPERM]

The entry to be unlinked is the mount point for a mounted file system. [EBUSY]

The entry to be unlinked is the last link to a pure procedure (shared text) file that is being executed. [ETXTBSY]

The directory entry to be unlinked is part of a read-only file system. [EROFS]

Path points outside the process's allocated address space. [EFAULT]

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

rm(1), close(2), link(2), open(2).

ASSEMBLER

```
moveq #10,D0
movl  path,A0
trap  #0
```

Carry bit set on failure and cleared on success.

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:

Dev is not the device number of a device containing a mounted file system. [EINVAL]

Buf points outside the process's allocated address space. [EFAULT]

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).

ASSEMBLER

```
moveq #57,D0
movl  buf,A0
movl  dev,D1
movl  #2,A1      |ustat
trap  #0
```

Carry bit set on failure and cleared on success.

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:

The named file does not exist. [ENOENT]

A component of the path prefix is not a directory. [ENOTDIR]

Search permission is denied by a component of the path prefix. [EACCES]

The effective user ID is not super-user and not the owner of the file and *times* is not NULL. [EPERM]

The effective user ID is not super-user and not the owner of the file and *times* is NULL and write access is denied. [EACCES]

The file system containing the file is mounted read-only. [EROFS]

Times is not NULL and points outside the process's allocated address space. [EFAULT]

Path points outside the process's allocated address space. [EFAULT]

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).

ASSEMBLER

```
moveq #30,D0
movl  path,A0
movl  times,D1
trap  #0
```

NAME

`uvar` — returns system-specific configuration information

SYNOPSIS

```
#include <sys/var.h>
```

```
uvar(v)  
struct var *v;
```

DESCRIPTION

Returns system-specific configuration information contained in the kernel. The information returned contains table sizes, mask words, and other system-specific information for programs such as *adb(1)*, *ld(1)*, and *ps(1)*.

Presently a maximum of 256 bytes of information is returned. This number is subject to change.

SEE ALSO

`/usr/include/sys/space.h`

ASSEMBLER

```
moveq #57,D0  
movl v,A0  
movw #33,A1  
trap #0
```

Carry bit is set if data could not be put into the address pointed to by `v`.

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 it receives a signal that is to be caught (see *signal(2)*), or until any one of the calling process's child processes stops in a trace mode (see *ptrace(2)*) or terminates. 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:

The calling process has no existing unwaited-for child processes. [ECHILD]

Stat_loc points to an illegal address. [EFAULT]

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)*, *pause(2)*, *signal(2)*.

WARNING

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

WAIT (2)

ASSEMBLER

```
moveq #7,D0
trap #0
bcs 2$
tstl stat_loc      | wait(0)?
beq 1$             | yes, return
movl stat_loc,A0
@ movl D1,A0@
```

Process ID in D0.
Status in D1.

Carry flag is set if there are no children not previously waited for.

WAIT(2)

WRITE (2)

WRITE (2)

NAME

write — write on a file

SYNOPSIS

```
int write (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.

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.

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

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

An attempt is made to write to a pipe that is not open for reading by any process. [EPIPE and SIGPIPE signal]

An attempt was made to write a file that exceeds the process's file size limit or the maximum file size. See *ulimit(2)*. [EFBIG]

Buf points outside the process's allocated address space. [EFAULT]

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), no partial writes will be permitted. Thus, the write will fail if a write of *nbyte* bytes would exceed a limit.

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.

SEE ALSO

creat(2), *dup(2)*, *lseek(2)*, *open(2)*, *pipe(2)*, *ulimit(2)*.

ASSEMBLER

```

moveq #4,D0
movl  fildes,A0
movl  buf,D1
movl  nbytes,A1
trap  #0

```

Carry bit set on failure and cleared on success.

The number of bytes written is returned in **D0**.

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.
- (3M) These functions constitute the Math Library, *libm*. 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>**.
- (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>**.
- (3X) Various specialized libraries. The files in which these libraries are found are given on the appropriate pages.

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 his own definition if he is not using `<stdio.h>`.

FILES

```

/lib/libc.a
/lib/libm.a

```

SEE ALSO

ar(1), *cc*(1), *fortran*(1), *ld*(1), *nm*(1), *intro*(2), *stdio*(3S).

DIAGNOSTICS

Functions in the Math Library (3M) may return the conventional values `0` or **HUGE** (the largest single-precision floating-point number) 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**.

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.

NAME

abort — generate an IOT fault

SYNOPSIS

`int abort ()`

DESCRIPTION

Abort 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

adb(1), *exit(2)*, *kill(2)*, *signal(2)*.

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).

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

cpp(1), *abort(3C)*.

NAME

atof — convert ASCII string to floating-point number

SYNOPSIS

```
double atof (nptr)  
char *nptr;
```

DESCRIPTION

Atof converts a character string pointed to by *nptr* to a double-precision floating-point number. The first unrecognized character ends the conversion. *Atof* recognizes an optional string of white-space characters (tabs and spaces), then an optional sign, then a string of digits optionally containing a decimal point, then an optional e or E followed by an optionally signed integer. If the string begins with an unrecognized character, *atof* returns the value zero.

DIAGNOSTICS

When the correct value would overflow, *atof* returns HUGE, and sets *errno* to ERANGE. Zero is returned on underflow.

SEE ALSO

scanf(3S), *strtol*(3C).

NAME

j0, *j1*, *jn*, *y0*, *y1*, *yn* — 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

J0 and *J1* return Bessel functions of *x* of the first kind of orders 0 and 1 respectively. *Jn* returns the Bessel function of *x* of the first kind of order *n*.

Y0 and *Y1* return the Bessel functions of *x* of the second kind of orders 0 and 1 respectively. *Yn* 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 *y0*, *y1* and *yn* to return the value **HUGE** and to set *errno* to **EDOM**. They also cause a message indicating **DOMAIN** error to be printed on the standard error output; the process will continue.

These error-handling procedures may be changed with the function *matherr*(3M).

SEE ALSO

matherr(3M).

NAME

blt, blt512 — block transfer data

SYNOPSIS

```
int blt(to,from,count)
char *to;
char *from;
int count;
```

```
int blt512(to,from,count)
char *to;
char *from;
int count;
```

DESCRIPTION

Blt does a fast copy of *count* bytes of data starting at address *from* to address *to*.

Blt512 does a fast copy of *count* number of consecutive 512 byte units starting at address *from* to address *to*.

NAME

bsearch — binary search

SYNOPSIS

```
char *bsearch ((char *) key, (char *) base, nel, width, compar)
unsigned nel, width;
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 the datum 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. *Width* is the width of an element in bytes; *sizeof (*key)* should be used. *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; accordingly, the first argument is to be considered less than, equal to, or greater than the second.

DIAGNOSTICS

A NULL pointer is returned if the key cannot be found in the table.

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

lsearch(3C), hsearch(3C), qsort(3C), tsearch(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)*.

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 lower-case letter, the result is the corresponding upper-case letter. If the argument of *tolower* represents an upper-case letter, the result is the corresponding lower-case letter. All other arguments in the domain are returned unchanged.

_toupper and *_tolower* are macros that accomplish the same thing as *toupper* and *tolower* but have restricted domains and are faster. *_toupper* requires a lower-case letter as its argument; its result is the corresponding upper-case letter. *_tolower* requires an upper-case letter as its argument; its result is the corresponding lower-case 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)*.

NAME

crypt, setkey, encrypt -- generate DES encryption

SYNOPSIS

```
char *crypt (key, salt)
char *key, *salt;
void setkey (key)
char *key;
void encrypt (block, edflag)
char *block;
int edflag;
```

DESCRIPTION

Crypt is the password encryption function. It is based on the NBS Data Encryption Standard (DES), with variations intended (among other things) to frustrate use of hardware implementations of the DES for 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 DES 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.

The *setkey* and *encrypt* entries provide (rather primitive) access to the actual DES algorithm. The argument of *setkey* is a character array of length 64 containing only the characters with numerical value 0 and 1. If 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. This is the key that will be used with the above mentioned algorithm to encrypt or decrypt the string *block* with the function *encrypt*.

The argument to the *encrypt* entry is a character array of length 64 containing only the characters with numerical value 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 DES algorithm using the key set by *setkey*. If *edflag* is zero, the argument is encrypted; if non-zero, it is decrypted.

SEE ALSO

login(1), passwd(1), getpass(3C), passwd(4).

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).

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 Savings 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 Savings 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 Savings Time conversion should be applied. The program knows

CTIME(3C)

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.

CTIME(3C)

CTYPE(3C)

CTYPE(3C)

NAME

isalpha, *isupper*, *islower*, *isdigit*, *isxdigit*, *isalnum*, *isspace*, *ispunct*, *isprint*, *isgraph*, *iscntrl*, *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 upper-case letter.

islower *c* is a lower-case 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.

iscntrl *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

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 of the owner of the current process. 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)*.

BUGS

Cuserid uses *getpwnam(3C)*; thus the results of a user's call to the latter will be obliterated by a subsequent call to the former.

The name *cuserid* is rather a misnomer.

NAME

dial — establish an out-going terminal line connection

SYNOPSIS

```
#include <dial.h>

int dial (call)
CALL *call;

void undial (fd)
int fd;
```

DESCRIPTION

Dial returns a file-descriptor for a terminal line open for read/write. The argument to *dial* is a CALL structure (defined in the <*dial.h*> header file.

When finished with the terminal line, the calling program must invoke *undial* to release the semaphore that has been set during the allocation of the terminal device.

The CALL typedef in the <*dial.h*> header file is:

```
typedef struct {
    struct termio *attr;    /* pointer to termio attribute struct */
    int          baud;     /* transmission data rate */
    int          speed;    /* 212A modem: low=300, high=1200 */
    char         *line;    /* device name for out-going line */
    char         *telno;   /* pointer to tel-no digits string */
    int          modem;    /* specify modem control for direct lines */
} CALL;
```

The CALL element *speed* is intended only for use with an outgoing dialed call, in which case its value should be either 300 or 1200 to identify the 113A modem, or the high or low speed setting on the 212A modem. The CALL element *baud* is for the desired transmission baud rate. For example, one might set *baud* to 110 and *speed* to 300 (or 1200).

If the desired terminal line is a direct line, a string pointer to its device-name should be placed in the *line* element in the CALL structure. Legal values for such terminal device names are kept in the *L-devices* file. In this case, the value of the *baud* element need not be specified as it will be determined from the *L-devices* file.

The *telno* element is for a pointer to a character string representing the telephone number to be dialed. The termination symbol will be supplied by the *dial* function, and should not be included in the *telno* string passed to *dial* in the CALL structure.

The CALL element *modem* is used to specify modem control for direct lines. This element should be non-zero if modem control is required. The CALL element *attr* is a pointer to a *termio* structure, as defined in the *termio.h* header file. A NULL value for this pointer element may be passed to the *dial* function, but if such a structure is included, the elements specified in it will be set for the outgoing terminal line before the connection is established. This is often important for certain attributes such as parity and baud-rate.

FILES

```
/usr/lib/uucp/L-devices
/usr/spool/uucp/LCK..tty-device
```

SEE ALSO

uucp(1C), alarm(2), read(2), write(2).
termio(7) in the *UniPlus+ Administrator's Manual*.

DIAGNOSTICS

On failure, a negative value indicating the reason for the failure will be returned. Mnemonics for these negative indices as listed here are defined in the `<dial.h>` header file.

```
INTRPT    -1    /* interrupt occurred */
D_HUNG    -2    /* dialer hung (no return from write) */
NO_ANS    -3    /* no answer within 10 seconds */
ILL_BD    -4    /* illegal baud-rate */
A_PROB    -5    /* acu problem (open() failure) */
L_PROB    -6    /* line problem (open() failure) */
NO_Ldv    -7    /* can't open LDEVS file */
DV_NT_A   -8    /* requested device not available */
DV_NT_K   -9    /* requested device not known */
NO_BD_A   -10   /* no device available at requested baud */
NO_BD_K   -11   /* no device known at requested baud */
```

WARNINGS

Including the `<dial.h>` header file automatically includes the `<termio.h>` header file.

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

An `alarm(2)` system call for 3600 seconds is made (and caught) within the `dial` module for the purpose of "touching" the `LCK..` file and constitutes the device allocation semaphore for the terminal device. Otherwise, `uucp(1C)` may simply delete the `LCK..` entry on its 90-minute clean-up rounds. The alarm may go off while the user program is in a `read(2)` or `write(2)` system call, causing an apparent error return. If the user program expects to be around for an hour or more, error returns from `reads` should be checked for (`errno == EINTR`), and the `read` possibly reissued.

NAME

`drand48`, `erand48`, `lrand48`, `nrand48`, `rand48`, `jrand48`, `srand48`, `seed48`, `lcong48` — generate uniformly distributed pseudo-random numbers

SYNOPSIS

```
double drand48 ( )
double erand48 (xsubi)
unsigned short xsubi[3];
long lrand48 ( )
long nrand48 (xsubi)
unsigned short xsubi[3];
long mrand48 ( )
long jrand48 (xsubi)
unsigned short xsubi[3];
void srand48 (seedval)
long seedval;
unsigned short *seed48 (seed16v)
unsigned short seed16v[3];
void lcong48 (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, 2³¹).

Functions `mrnd48` and `jrnd48` return signed long integers uniformly distributed over the interval [-2³¹, 2³¹).

Functions `srand48`, `seed48` and `lcong48` are initialization entry points, one of which should be invoked before either `drand48`, `lrand48` or `mrnd48` is called. (Although it is not recommended practice, constant default initializer values will be supplied automatically if `drand48`, `lrand48` or `mrnd48` is called without a prior call to an initialization entry point.) Functions `erand48`, `nrand48` and `jrnd48` 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_n , according to the linear congruential formula

$$X_{n+1} = (aX_n + c) \bmod m \quad n \geq 0.$$

The parameter $m = 2^{48}$; hence 48-bit integer arithmetic is performed. Unless `lcong48` 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`, `mrnd48` or `jrnd48` 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 (left-most) 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 *srnd48* 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 *srnd48* or *seed48* will restore the “standard” multiplier and addend values, a and c , specified on the previous page.

NOTES

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).

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;
char *buf;
```

DESCRIPTION

Ecvt converts *value* to a null-terminated string of *ndigit* digits and returns a pointer thereto. 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 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 return values point to static data whose content is overwritten by each call.

NAME

end, *etext*, *edata* — last locations in program

SYNOPSIS

```
extern end;  
extern 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.

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. Thus, the current value of the program break should be determined by *sbrk(0)* (see *brk(2)*).

These symbols are accessible from assembly language if it is remembered that they should be prefixed by *_*.

SEE ALSO

brk(2), *malloc(3C)*.

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 . The values of x and y may not both be zero. If x is non-positive, y must be an integer.

Sqrt returns the square root of x . The value of x may not be negative.

DIAGNOSTICS

Exp returns HUGE when the correct value would overflow, and sets *errno* to ERANGE.

Log and *log10* return 0 and set *errno* to EDOM when x is non-positive. An error message is printed on the standard error output.

Pow returns 0 and sets *errno* to EDOM when x is non-positive and y is not an integer, or when x and y are both zero. In these cases a message indicating DOMAIN error is printed on the standard error output. When the correct value for *pow* would overflow, *pow* returns HUGE 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

intro(2), *hypot*(3M), *matherr*(3M), *sinh*(3M).

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

feof, feof, clearerr, fileno — stream status inquiries

SYNOPSIS

```
#include <stdio.h>
int feof (stream) FILE *stream;
int ferror (stream) FILE *stream;
void clearerr (stream) FILE *stream;
int fileno(stream) FILE *stream;
```

DESCRIPTION

Feof returns non-zero when EOF has previously been detected reading the named input *stream*, otherwise zero.

Ferror returns non-zero when an I/O error has previously occurred reading from or writing to the named *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)*.

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 number f with the same sign as x , such that $x = iy + f$ for some integer i , and $|f| < |y|$. *Fmod* will thus return x if y is zero.

Fabs returns $|x|$.

SEE ALSO

abs(3C).

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 obtained from *open*, *dup*, *creat*, or *pipe(2)*, which will open files but not return pointers to a FILE structure *stream* which 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

`open(2)`, `fclose(3S)`.

DIAGNOSTICS

`Fopen` and `freopen` return a NULL pointer on failure.

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 beginning at `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 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 variable `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 `nitems` is non-positive, no characters are read or written and 0 is returned by both `fread` and `fwrite`.

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 $x * 2^n$, 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*.

Ldexp returns the quantity $value * 2^{exp}$.

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, HUGE is returned and *errno* is set to ERANGE.

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)*.

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 a 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.

NAME

gamma – log gamma function

SYNOPSIS

```
#include <math.h>
extern int signgam;
double gamma (x)
double x;
```

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 Γ :

```
if ((y = gamma(x)) > LOGHUGE)
    error();
y = signgam * exp(y);
```

where LOGHUGE is the least value that causes *exp(3M)* to return a range error.

DIAGNOSTICS

For non-negative integer arguments HUGE is returned, and *errno* is set to EDOM. A message indicating DOMAIN 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)*.

NAME

getc, *getchar*, *fgetc*, *getw* — get character or word from 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*. It also moves the file pointer, if defined, ahead one character in *stream*. *Getc* is a macro and so cannot be used if a function is necessary; for example one cannot have a function pointer point to it.

Getchar returns the next character from the standard input stream, *stdin*. As in the case of *getc*, *getchar* is a macro.

Fgetc performs the same function as *getc*, but is a genuine function. *Fgetc* runs more slowly than *getc*, but takes less space per invocation.

Getw returns the next word (32-bit integer on a 68000) from the named input *stream*. It returns the constant EOF upon end-of-file or error, but as that is a valid integer value, *feof* and *ferror* (3S) should be used to check the success of *getw*. *Getw* increments the associated file pointer, if defined, to point to the next word. *Getw* assumes no special alignment in the file.

SEE ALSO

fclose(3S), *ferror*(3S), *fopen*(3S), *fread*(3S), *gets*(3S), *putc*(3S), *scanf*(3S).

DIAGNOSTICS

These functions return the integer constant EOF at end-of-file or upon an error.

BUGS

Because it is implemented as a macro, *getc* treats incorrectly a *stream* argument with side effects. In particular, *getc(*f++)* doesn't 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

pwd(1), *malloc*(3C), *popen*(3S).

DIAGNOSTICS

Returns NULL with *errno* set if *size* is not large enough, or if an error occurs in a lower-level function.

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

environ(5).

NAME

getgrent, getgrgid, getgrnam, setgrent, endgrent — 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 ( )
```

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.

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.

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* 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;
char *optstring;
extern char *optarg;
extern int optind;
```

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*.

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.

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 int optind;
    extern char *optarg;
    :
    :
    while ((c = getopt (argc, argv, "abf:o:")) != EOF)
        switch (c) {
            case 'a':
                if (bflg)
                    errflg++;
                else
                    aflg++;
                break;
            case 'b':
                if (aflg)
```

```

                errflg++;
            else
                bproc( );
            break;
        case 'f':
            ifile = optarg;
            break;
        case 'o':
            ofile = optarg;
            bufsiza = 512;
            break;
        case '?':
            errflg++;
    }
    if (errflg) {
        fprintf( stderr, "usage: . . . ");
        exit (2);
    }
    for ( ; optind < argc; optind++) {
        if (access (argv[optind], 4)) {
            :
        }
    }
}

```

SEE ALSO
getopt(1).

NAME

getpass — read a password

SYNOPSIS

```

char *getpass (prompt)
char *prompt;

```

DESCRIPTION

Getpass reads up to a newline 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.

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. The line is null-terminated. *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 the size of programs, not otherwise using standard I/O, more than might be expected.

NAME

getpwent, getpwuid, getpwnam, setpwent, endpwent — 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 ( )
```

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;
};

struct comment {
    char      *c_dept;
    char      *c_name;
    char      *c_acct;
    char      *c_bin;
};
```

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.

FILES

/etc/passwd

SEE ALSO

cuserid(3S), 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. Also see *cuserid(3S)*.

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.

NOTE

Gets deletes the new-line ending its input, but *fgets* keeps it.

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, inxx) */
    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 doesn't 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 isn't 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.

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 <i>remainder modulo table size</i> 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 <i>hcompare</i> and should behave in a manner similar to <i>strcmp</i> (see <i>string(3C)</i>). |
| 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. |

HSEARCH (3C)**HSEARCH (3C)**

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.

SEE ALSO

bsearch(3C), lsearch(3C), 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.

BUGS

Only one hash search table may be active at any given time.

HYPOT (3M)**HYPOT (3M)****NAME**

hypot — Euclidean distance function

SYNOPSIS

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

```
double hypot (x, y)
```

```
double x, y;
```

DESCRIPTION

Hypot returns

```
sqrt(x * x + y * 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

l3tol, *ltol3* — convert between 3-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.

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

env(1), **login(1)**, **profile(4)**, **environ(5)**.

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

lsearch — linear search and update

SYNOPSIS

```
char *lsearch ((char *)key, (char *)base, nelp, width, compar)
unsigned *nelp, width;
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. *Nelp* points to an integer containing the current number of elements in the table. The integer is incremented if the datum is added to the table. *Width* is the width of an element in bytes; *sizeof (*key)* should be used. *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.

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), hsearch(3C), tsearch(3C)
The Art of Computer Programming, Volume 1, Sorting and Searching by Donald Knuth.

BUGS

Undefined results can occur if there is not enough room in the table to add a new item.

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;

cfree (ptr, nelem, elsize)
    char *ptr;
    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.

The arguments to *cfree* are the pointer to a block previously allocated by *calloc* plus the parameters to *calloc*.

Each of the allocation routines returns a pointer to space suitably aligned (after possible pointer coercion) for storage of any type of object.

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.

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. A pointer to the exception structure *x* will be passed to the user-supplied *matherr* function when an error occurs. 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	domain error
SING	singularity
OVERFLOW	overflow
UNDERFLOW	underflow
TLOSS	total loss of significance
PLOSS	partial loss of significance

The element *name* points to a string containing the name of the function that had the error. The variables *arg1* and *arg2* are the arguments to the function that had the error. *Retval* is a double that is returned by the function having the error. If it supplies a return value, the user's *matherr* must return non-zero. If the default error value is to be returned, the user's *matherr* must return 0.

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 non-zero and the program continues.

EXAMPLE

```
matherr(x)
register struct exception *x;
{
    switch (x->type) {
    case DOMAIN:
    case SING: /* print message and abort */
        fprintf(stderr, "domain error in %s\n", x->name);
        abort();
    case OVERFLOW:
        if (!strcmp("exp", x->name)) {
            /* if exp, print message, return the argument */
            fprintf(stderr, "exp of %f\n", x->arg1);
```

```

        x->retval = x->arg1;
    } else if (!strcmp("sinh", x->name)) {
        /* if sinh, set errno, return 0 */
        errno = ERANGE;
        x->retval = 0;
    } else
        /* otherwise, return HUGE */
        x->retval = HUGE;
    break;
case UNDERFLOW:
    return (0); /* execute default procedure */
case TLOSS:
case PLOSS:
    /* print message and return 0 */
    fprintf(stderr, "loss of significance in %s\n", x->name);
    x->retval = 0;
    break;
}
return (1);
}
    
```

DEFAULT ERROR HANDLING PROCEDURES						
	Types of Errors					
	DOMAIN	SING	OVERFLOW	UNDERFLOW	TLOSS	PLOSS
BESSEL: y0, y1, yn (neg. no.)	- M, -H	-	H	0	-	*
EXP:	-	-	H	0	-	-
POW: (neg.)**(non- int.), 0**0	M, 0	-	H	0	-	-
LOG: log(0): log(neg.):	- M, -H	M, -H	-	-	-	-
SQRT:	M, 0	-	-	-	-	-
GAMMA:	-	M, H	-	-	-	-
HYPOT:	-	-	H	-	-	-
SINH, COSH:	-	-	H	-	-	-
SIN, COS:	-	-	-	-	M, 0	M, *
TAN:	-	-	H	-	0	*
ACOS, ASIN:	M, 0	-	-	-	-	-

ABBREVIATIONS	
*	As much as possible of the value is returned.
M	Message is printed.
H	HUGE is returned.
-H	-HUGE is returned.
0	0 is returned.

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 efficiently 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.

Character movement is performed differently in different implementations. Thus overlapping moves may yield surprises.

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

```
void monitor (lowpc, highpc, buffer, bufsize, nfunc)
int (*lowpc) ( ), (*highpc) ( );
short *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* short integers. *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 etext;
...
monitor ((int (*)())2, etext, buf, bufsize, nfunc);
```

Etext 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 ((int (*)())NULL, 0, 0, 0, 0);
```

Prof(1) can then be used to examine the results.

FILES

`mon.out`

SEE ALSO

cc(1), *prof(1)*, *profil(2)*, *end(3C)*.

NAME

nlist — get entries from name list

SYNOPSIS

```
#include <a.out.h>
int nlist (file-name, nl)
char *file-name;
struct nlist *nll l;
```

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. 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 subroutine 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.

SEE ALSO

a.out(4).

DIAGNOSTICS

All type entries are set to 0 if the file cannot be read or if it doesn't contain a valid name list.

Nlist returns -1 upon error; otherwise it returns 0.

NAME

perror, errno, sys_errlist, sys_nerr — system error messages

SYNOPSIS

```
void perror (s)
char *s;
extern int errno;
extern char *sys_errlist[];
extern int sys_nerr;
```

DESCRIPTION

Perror 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).

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 tplot(1G) 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

PLOT (3X)

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.

SEE ALSO

tplot(1G), plot(4).

PLOT (3X)**POPEN (3S)****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 *fflush*; 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 (see below) has been given) to the field width;

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** specifying that a following **d**, **o**, **u**, **x**, or **X** conversion character applies to a long integer *arg*.

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.

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* (*X*) conversion, a non-zero result will have *0x* (*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. 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, 6 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, 6 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 ($\backslash 0$) 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. If the string pointer *arg* has the value zero, the result is

- undefined. A *null* 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 *sprintf* are printed as if *putc* (3S) had been called.

EXAMPLE

```
printf("%s, %s %d, %.2d:%.2d", weekday, month, day, hour, min);
```

prints a date and time in the form "Sunday, July 3, 10:02", where *weekday* and *month* are pointers to null-terminated strings.

```
printf("pi = %.5f", 4*atan(1.0));
```

prints π to 5 decimal places.

SEE ALSO

ecvt(3C), *putc*(3S), *scanf*(3S), *stdio*(3S).

NAME

putc, putchar, fputc, putw — put character or word on a stream

SYNOPSIS

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

```
int putc (c, stream)
```

```
char c;
```

```
FILE *stream;
```

```
int putchar (c)
```

```
char c;
```

```
int fputc (c, stream)
```

```
char 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 takes less space per invocation.

Putw writes the word (32-bit integer on the 68000) *w* to the output *stream* (at the position at which the file pointer, if defined, is pointing). *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)* 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, these functions 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. Because EOF is a valid integer, *ferror(3S)* should be used to detect *putw* errors.

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. For this reason the use of *putw* should be avoided.

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*), *putpwuid* 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.

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), putchar(3S).

NOTES

Puts appends a new-line character while *fputs* does not.

NAME

qsort — quicker sort

SYNOPSIS

```
void qsort ((char *) base, nel, width, compar)
unsigned int nel, width;
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. *Width* is the width of an element in bytes; *sizeof (*base)* should be used. *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 according as the first argument is to be considered 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.

EXAMPLE

```
struct entry {
    char *name;
    int flags;
};
main()
{
    struct entry hp[100];
    int entcmp();
    int i, count;
    for (i = 0; i < (count = 100); i++) {
        /* fill the structure with the name and flags */
        :
    }
    qsort( (char *) hp, count, sizeof (hp[0]), entcmp);
}
entcmp(ep,ep2)
struct entry *ep, *ep2;
{
    return (strcmp(ep->name, ep2->name));
}
```

will sort a set of names with associated flags in ASCII order.

SEE ALSO

sort(1), bsearch(3C), lsearch(3C), string(3C).

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).

NAME

regcmp, regex — compile and execute regular expression

SYNOPSIS

```
char *regcmp(string1 [, string2, ...], 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 the 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 last or first 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. *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. {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 present, 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 parenthesis. For example, (a*(cb+)*)\$0.

By necessity, all the above defined symbols are special. They must, therefore, be escaped to be used as themselves.

EXAMPLE

```
char *cursor, *newcursor, *ptr;
...
newcursor = regex((ptr = regcmp("^\n", 0)), cursor);
free(ptr);
```

matches a leading new-line in the subject string pointed at by cursor.

```
char ret0[9];
char *newcursor, *name;
...
name = regcmp("[A-Za-z][A-Za-z0-9_]{0,7})$0", 0);
newcursor = regex(name, "123Testing321", ret0);
```

matches 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*.

```
#include "file.i"
char *string, *newcursor;
...
newcursor = regex(name, string);
```

applies a precompiled regular expression in *file.i* (see *regcmp(1)*) against *string*.

This routine is kept in */lib/libPW.a*.

SEE ALSO

ed(1), *regcmp(1)*, *malloc(3C)*.

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 */
...
malloc(n) {
    static int rebuf[256];
    return rebuf;
}
```

NAME

rhost, *raddr* — look up internet hosts by name or address

SYNOPSIS

```
iaddr = rhost(aname)
long iaddr;
char **aname;
```

```
name = raddr(iaddr)
long iaddr;
```

DESCRIPTION

Rhost is given a pointer to a name for an Internet host and returns the 32 bit internet address in network byte order suitable for direct use in a **sockaddr_in** internet address as **sockaddr_in.sin_addr.s_addr**. If the host name is not known then *rhost* returns -1. If the host name is known then ***aname** is changed to point to the standard name of the specified host, which is the first name given in its entry in */etc/hosts*. The return value has been saved with *malloc* and is not destroyed on subsequent calls.

Raddr performs a similar function, but takes an Internet address, and looks up the name.

FILES

/etc/hosts

SEE ALSO

remsh(1), *rlogin(1)*, *socket(2)*.

BUGS

A more general data base or server is needed.

This interface is provisional and may be changed in future releases.

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 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.

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 should be 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 optionally signed 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.
- l 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.

The conversion characters `d`, `u`, `o`, and `x` may be preceded by `l` or `h` 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.

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.

EXAMPLE

The call:

```
int i; float x; char name[50];
scanf ("%d%f%s", &i, &x, name);
```

with the input line:

```
25 54.32E-1 thompson
```

assigns 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];
scanf ("%2d%f%*d %[0-9]", &i, &x, name);
```

with input:

```
56789 0123 56a72
```

assigns 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

atof(3C), *getc*(3S), printf(3S), 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 — assign buffering to a stream

SYNOPSIS

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

```
void setbuf (stream, buf)
```

```
FILE *stream;
```

```
char *buf;
```

DESCRIPTION

Setbuf is used after a stream has been opened but before it is read or written. It causes the character array pointed to by *buf* to be used instead of an automatically allocated buffer. If *buf* is a NULL character 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];
```

A buffer is normally obtained from *malloc(3C)* at the time of the first *getc* or *putc(3S)* on the file, except that the standard error stream *stderr* is normally not buffered.

Output streams directed to terminals are always line-buffered unless they are unbuffered.

SEE ALSO

fopen(3S), *getc(3S)*, *malloc(3C)*, *putc(3S)*.

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.

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 have values as of the time *longjmp* was called.

SEE ALSO

signal(2).

WARNING

If *longjmp* is called when *env* was never primed by a call to *setjmp*, or when the last such call is in a function which has since returned, absolute chaos is guaranteed.

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 real argument.

DIAGNOSTICS

Sinh and *cosh* return HUGE 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).

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, and 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).

NAME

sputl, *sgetl* — access long numeric data in a machine independent fashion.

SYNOPSIS

```
sputl ( value, buffer )  
long value;  
char *buffer;  
  
long sgetl ( buffer )  
char *buffer;
```

DESCRIPTION

Sputl(3X) will take the 4 bytes of the long *value* and place them in memory starting at the address pointed to by *buffer*. The ordering of the bytes is the same across all machines. *Sgetl* will retrieve the 4 bytes in memory starting at the address pointed to by *buffer* and return the long value in the byte ordering of the host machine.

The usage of *sputl*(3X) and *sgetl* in combination provides a machine independent way of storing long numeric data in an ASCII file. The numeric data stored in the portable archive file format (see *ar*(4)) is written and read into/from buffers with *sputl*(3X) and *sgetl* respectively.

A program which uses these functions must be loaded with the object file access routine library *libld.a*.

SEE ALSO

ar(4).

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.

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*, *putchar*, and the higher-level routines *fgetc*, *fgets*, *sprintf*, *fputc*, *fputs*, *fread*, *fscanf*, *fwrite*, *gets*, *getw*, *printf*, *puts*, *putw*, and *scanf* all 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).

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*, *ferror*, *clearerr*, and *fileno*.

SEE ALSO

open(2), *close*(2), *lseek*(2), *pipe*(2), *read*(2), *write*(2), *ctermid*(3S), *cuserid*(3S), *fclose*(3S), *ferror*(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.

NAME

stdipc — 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 *strchr (s, c)
char *s, c;

char *strrchr (s, c)
char *s, 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.

STRING (3C)

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* or more. Each function returns *s1*.

Strlen returns the number of characters in *s*, not including the terminating null character.

Strchr (*strchr*) 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 on 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 uses native character comparison.

All string movement is performed character by character starting at the left. Thus overlapping moves toward the left will work as expected, but overlapping moves to the right may yield surprises.

STRING (3C)**STRTOL (3C)****STRTOL (3C)****NAME**

strtol, *atol*, *atoi* — convert string to integer

SYNOPSIS

long strtol (str, ptr, base)

char *str;
char **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 *str*. The string is scanned up to the first character inconsistent with the base. Leading “white-space” characters are ignored.

If the value of *ptr* is not (char **)NULL, a pointer to the character terminating the scan is returned in **ptr*. If no integer can be formed, **ptr* 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 thus: 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

atof(3C), *scanf*(3S).

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.

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

sh(1), *exec*(2).

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 `-1` and sets *errno*.

NAME

tgetent, tgetnum, tgetflag, tgetstr, tgoto, tputs — terminal independent operation routines

SYNOPSIS

```
char PC;
char *BC;
char *UP;
short ospeed;

tgetent(bp, name)
char *bp, *name;

tgetnum(id)
char *id;

tgetflag(id)
char *id;

char *
tgetstr(id, area)
char *id, **area;

char *
tgoto(cm, destcol, destline)
char *cm;

tputs(cp, affcnt, outc)
register char *cp;
int affcnt;
int (*outc)();
```

DESCRIPTION

These functions extract and use capabilities from the terminal capability data base *termcap*(5). Note that these are low level routines.

Tgetent extracts the entry for terminal *name* into the buffer at *bp*. *Bp* should be a character buffer of size 1024 and must be retained through all subsequent calls to *tgetnum*, *tgetflag*, and *tgetstr*. *Tgetent* returns -1 if it cannot open the *termcap* file, 0 if the terminal name given does not have an entry, and 1 if all goes well. It will look in the environment for a TERMCAP variable. If found, and the value does not begin with a slash, and the terminal type *name* is the same as the environment string TERM, the TERMCAP string is used instead of reading the *termcap* file. If it does begin with a slash, the string is used as a path name rather than */etc/termcap*. This can speed up entry into programs that call *tgetent*, as well as to help debug new terminal descriptions or to make one for your terminal if you can't write the file */etc/termcap*.

Tgetnum gets the numeric value of capability *id*, returning -1 if it is not given for the terminal. *Tgetflag* returns 1 if the specified capability is present in the terminal's entry, 0 if it is not. *Tgetstr* gets the string value of capability *id*, placing it in the buffer at *area*, advancing the *area* pointer. It decodes the abbreviations for this field described in *termcap*(5), except for cursor addressing and padding information.

Tgoto returns a cursor addressing string decoded from *cm* to go to column *destcol* in line *destline*. It uses the external variables UP (from the *up* capability) and BC (if *bc* is given rather than *bs*) if necessary to avoid placing *\n*, *^D* or *^@* in the returned string. (Programs which call *tgoto* should be

TERMCAP(3)**TERMCAP(3)**

sure to turn off the XTABS bit(s), since *tgoto* may now output a tab. Note that programs using *termcap* should in general turn off XTABS anyway since some terminals use control-I for other functions, such as nondestructive space.) If a % sequence is given which is not understood, then *tgoto* returns OOPS.

Tputs decodes the leading padding information of the string *cp*; *affcnt* gives the number of lines affected by the operation, or 1 if this is not applicable, *outc* is a routine which is called with each character in turn. The external variable *ospeed* should contain the output speed of the terminal as encoded by *stty*(2). The external variable PC should contain a pad character to be used (from the *pc* capability) if a null (^@) is inappropriate.

FILES

/usr/lib/libtermcap.a	termcap library
/etc/termcap	data base

SEE ALSO

ex(1), termcap(5).

AUTHOR

William Joy

TMPFILE(3S)**TMPFILE(3S)****NAME**

tmpfile — create a temporary file

SYNOPSIS

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

```
FILE *tmpfile ()
```

DESCRIPTION

Tmpfile creates a temporary file and returns a corresponding FILE pointer. The file will automatically be deleted when the process using it terminates. The file is opened for update.

SEE ALSO

creat(2), unlink(2), fopen(3S), mktemp(3C), 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-name 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 path-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 path-name for an appropriate directory, the path-name defined as **P_tmpdir** in the `<stdio.h>` header file is used. If that path-name 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 a path-name for 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* 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* or *creat* 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

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 x, y;
```

DESCRIPTION

Sin, *cos* and *tan* return respectively the sine, cosine and tangent of their argument, which is 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 π .

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 π , 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 0 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, a PLOSS error is generated but no message is printed. In both cases, *errno* is set to ERANGE.

Tan returns HUGE for an argument which is near an odd multiple of $\pi/2$ when the correct value would overflow, and sets *errno* to ERANGE.

Arguments of magnitude greater than 1.0 cause *asin* and *acos* to return 0 and to set *errno* 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).

NAME

tsearch, *tdelete*, *twalk* — manage binary search trees

SYNOPSIS

```
#include <search.h>
char *tsearch ((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 is a binary tree search routine generalized from Knuth (6.2.2) Algorithm T. It returns a pointer into a tree indicating where a datum may be found. If the datum does not occur, it is added at an appropriate point in the tree. *Key* points to the datum to be sought in the tree. *Rootp* points to a variable that points to the root of the tree. A NULL pointer value for the variable denotes an empty tree; in this case, the variable will be set to point to the datum at the root of the new tree. *Compar* is the name of the comparison function. It 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 according as the first argument is to be considered less than, equal to, or greater than the second.

Tdelete deletes a node from a binary search tree. It is generalized from Knuth (6.2.2) algorithm D. 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.

NOTES

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. 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.

Warning: the *root* argument to *twalk* is one level of indirection less than the *rootp* arguments to *tsearch* and *tdelete*.

DIAGNOSTICS

A NULL pointer is returned by *tsearch* if there is not enough space available to create a new node.

TSEARCH(3C)

TSEARCH(3C)

A NULL pointer is returned by *tsearch* and *tdelete* if *rootp* is NULL on entry.

SEE ALSO

bsearch(3C), *hsearch(3C)*, *lsearch(3C)*.

BUGS

Awful things can happen if the calling function alters the pointer to the root.

TTYNAME(3C)

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), *ttynam(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.

NAME

`ungetc` — push character back into input stream

SYNOPSIS

```
#include <stdio.h>
int ungetc (c, stream)
char 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* call on that *stream*. *Ungetc* returns *c*, and leaves the file *stream* unchanged.

One character of pushback is guaranteed provided something has been read from the stream and the stream is actually buffered.

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

In order that *ungetc* perform correctly, a read statement must have been performed prior to the call of the *ungetc* function. *Ungetc* returns EOF if it can't insert the character. In the case that *stream* is *stdin*, *ungetc* will allow exactly one character to be pushed back onto the buffer without a previous read statement.

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 *UniPlus⁺ Administrator's Manual*.

NAME

a.out — assembler and link editor output

SYNOPSIS

#include <a.out.h>

DESCRIPTION

A.out is the output file of the assembler *as(1)* and the link loader *ld(1)*. *Ld(1)* makes *a.out* executable if there were no errors and no unresolved external references. Layout information as given in the include file for the 68000 is:

```

/ *
 * Layout of a.out file:
 *
 * header of 8 longs   magic number 405, 407, 410, 411
 *                   text size
 *                   data size
 *                   bss size
 *                   symbol table size
 *                   text relocation size
 *                   data relocation size
 *                   entry point
 *
 * header:            0
 * text:              32
 * data:              32+textsize
 * symbol table:     32+textsize+datasize
 * text relocation:  32+textsize+datasize+symsize
 * data relocation:  32+textsize+datasize+symsize+rtextsize
 *
 */

/* various parameters */
#define SYMLENGTH 50 /* maximum length of a symbol */

/* types of files */
#define ARCMAGIC 0177545 /* ar files */
#define FMAGIC 0407 /* standard executable */
#define NMAGIC 0410 /* shared text executable */

/* symbol types */
#define EXTERN 040 /* external */
#define UNDEF 00 /* undefin d */
#define ABS 01 /* absolute */
#define TEXT 02 /* text */
#define DATA 03 /* data */
#define BSS 04 /* bss */
#define COMM 05 /* internal use only */
#define REG 06 /* register name */

/* relocation regions */
#define RTEXT 00
#define RDATA 01
#define RBSS 02
#define REXT 03

```

```

/* relocation sizes */
#define RBYTE 00
#define RWORD 01
#define RLONG 02

/* macros which define various positions in file based on a bhdr, filhdr */
#define TEXTPOS ((long) sizeof(filhdr))
#define DATAPOS (TEXTPOS + filhdr.tsize)
#define SYMPOS (DATAPOS + filhdr.dsize)
#define RTEXTPOS (SYMPOS + filhdr.ssize)
#define RDATAPOS (RTEXTPOS + filhdr.rtsize)
#define ENDPOS (RDATAPOS + filhdr.rdsiz)

/* header of a.out files */
struct bhdr {
    long    fmagic;
    long    tsize;
    long    dsize;
    long    bsize;
    long    ssize;
    long    rtsize;
    long    rdsiz;
    long    entry;
};

/* symbol management */
struct sym {
    char    stype;        /* symbol type */
    char    sympad;      /* pad to short align */
    long    svalue;      /* value */
};

/* relocation commands */
struct reloc {
    unsigned rsegment:2; /* RTEXT, RDATA, RBSS, or REXTERN */
    unsigned rsize:2;    /* RBYTE, RWORD, or RLONG */
    unsigned rdisp:1;    /* 1 => a displacement */
    unsigned relpad1:3;  /* pad 1 */
    char     relpad2;    /* pad 2 */
    short    rsymbol;    /* id of the symbol of external relocations */
    long     rpos;       /* position of relocation in segment */
};

/* symbol table entry */
struct nlist {
    char    n_name[8];   /* symbol name */
    int     n_type;      /* type flag */
    unsigned n_value;    /* value */
};

```

```

/* values for type flag */
#define N_UNDF 0        /* undefined */
#define N_ABS 01       /* absolute */
#define N_TEXT 02      /* text symbol */
#define N_DATA 03      /* data symbol */
#define N_BSS 04       /* bss symbol */
#define N_TYPE 037
#define N_REG 024      /* register name */
#define N_FN 037       /* file name symbol */
#define N_EXT 040      /* external bit, or'ed in */
#define FORMAT "%06o" /* to print a value */

```

The file has four sections: a header, the program and data text, a symbol table, and relocation information. The last two may be empty if the program was loaded with the `-s` option of `ld` or if the symbols and relocation have been removed by `strip(1)`.

In the header the sizes of each section are given in bytes, but are even. The size of the header is not included in any of the other sizes.

When an `a.out` file is loaded into core for execution, three logical segments are set up: the text segment, the data segment (with uninitialized data, which starts off as all 0, following initialized data), and a stack. The text segment begins at the user program start address in the core image; the header is not loaded. If the magic number in the header is `FMAGIC`, it indicates that the text segment is not to be write-protected and shared, so the data segment is immediately contiguous with the text segment. If the magic number is `NMAGIC`, the data segment begins at the next segment boundary following the text segment, and the text segment is not writable by the program; if other processes are executing the same file, they will share the text segment.

The stack will occupy the highest possible user program locations in the core image and will grow downwards. The stack is automatically extended as required. The data segment is only extended as requested by `brk(2)`.

The start of the text segment in the file is `32(10)`; the start of the data segment is `32+St` (the size of the text) the start of the relocation information is `32+St+Sd`; the start of the symbol table is `32+2(St+Sd)` if the relocation information is present, `32+St+Sd` if not.

The layout of a symbol table entry and the principal flag values that distinguish symbol types are given in the include file.

If a symbol's type is undefined external, and the value field is non-zero, the symbol is interpreted by the loader `ld` as the name of a common region whose size is indicated by the value of the symbol.

The value of a word in the text or data portions which is not a reference to an undefined external symbol is exactly that value which will appear in core when the file is executed. If a word in the text or data portion involves a reference to an undefined external symbol, as indicated by the relocation information for that word, then the value of the word as stored in the file is an offset from the associated external symbol. When the file is processed by the link editor and the external symbol becomes defined, the value of the symbol will be added into the word in the file.

If relocation information is present, it will appear in the form of the structure shown above.

SEE ALSO

as(1), ld(1), nm(1)

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;          /* Accounting user ID */
    ushort ac_gid;          /* Accounting group ID */
    dev_t  ac_tty;          /* control typewriter */
    time_t ac_btime;        /* Beginning time */
    comp_t ac_ftime;        /* 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 $ac_mem / (ac_stime + ac_ftime)$ can be viewed as an approximation to the mean process size, as modified by text-sharing.

The structure `tacct`, 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(1M)`, `acctcom(1)`, `acct(2)`.

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.

JAME

`altblk` — alternate block information for bad block handling

SYNOPSIS

```
#include <altblk.h>
```

DESCRIPTION

`Altblk` is the data structure used by `badblk(1M)` to handle bad blocks for disk drives that support soft sector bad block remapping.

The layout of this structure is as follows:

```
#define  MAXALT      50      /* max alternate disk blocks */
#define  ALTMAGIC    0xDBDF /* bad block information is valid flag */
/*
 * structure for alternate block mapping
 */
struct a_map {
    long a_altbk; /* bad block */
    long a_index; /* relative bad block index */
};
/*
 * disk header block format for alternate block mapping
 */
struct altblk {
    char a_fill[BSIZE-sizeof(struct a_map)-4*sizeof(long)];
    struct a_map a_map[1]; /* mapping */
    long a_magic; /* verification code (ALTMAGIC) */
    long a_count; /* bad block count */
    long a_nicbad; /* max number of bad blocks */
    long a_maxalt; /* max alt block used so far */
};
```

This structure describes the upper portion of block 0 of each physical disk. The array `a_map` is inverted (i.e., it is indexed backwards). The specific fields in `altblk` are:

`a_maxalt` — the next usable block in bad block area relative to the start of the bad block area

`a_nicbad` — the maximum number of elements in the `a_map` structure

`a_count` — the number of bad blocks currently remapped on the disk

`a_magic` — a magic number for verification

`a_map` — bad block remap information

SEE ALSO

`badblk(1M)`

NAME

ar — archive (library) file format

SYNOPSIS

```
#include <ar.h>
```

DESCRIPTION

The archive command *ar* is used to combine several files into one. Archives are used mainly as libraries to be searched by the link-editor *ld*.

A file produced by *ar* has a magic number at the start, followed by the constituent files, each preceded by a file header. The magic number and header layout as described in the include file are:

```
#define ARFMAG 0177545
```

```
struct ar_hdr {
    char    ar_name[14];
    long    ar_date;
    short   ar_uid;
    short   ar_gid;
    short   ar_mode;
    long    ar_size;
};
```

The "ar_fm_ag" field contains the 32-bit number ARFMAG to help verify the presence of a header. The name is a blank padded string. The other fields are left-adjusted, blank-padded numbers. They are decimal except for "ar_mode", which is octal. The date is the modification date of the file at the time of its insertion into the archive.

Each file begins on an even (0 mod 2) boundary; a new-line is inserted between files if necessary. Nevertheless the size given reflects the actual size of the file exclusive of padding.

There is no provision for empty areas in an archive file.

SEE ALSO

ar(1), ld(1), nm(1)

BUGS

File names lose trailing blanks. Most software dealing with archives takes even an included blank as a name terminator.

CHECKLIST (4)

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.

FILES

/etc/checklist

SEE ALSO

fsck(1M).

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 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, or separated from data space, 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**.

SEE ALSO

setuid(2), signal(2).

NAME

cpio — format of cpio archive

DESCRIPTION

The *header* structure, when the `-c` option of *cpio*(1) is not used, is:

```
struct {
    short    h_magic,
            h_dev;
    ushort   h_ino,
            h_mode,
            h_uid,
            h_gid;
    short    h_nlink,
            h_rdev,
            h_mtime[2],
            h_namesize,
            h_filesize[2];
    char     h_name[h_namesize rounded to word];
} Hdr;
```

When the `-c` option is used, the *header* information is described by:

```
sscanf(Chdr,"%6o%6o%6o%6o%6o%6o%6o%6o%11lo%6o%11lo%s",
        &Hdr.h_magic, &Hdr.h_dev, &Hdr.h_ino, &Hdr.h_mode,
        &Hdr.h_uid, &Hdr.h_gid, &Hdr.h_nlink, &Hdr.h_rdev,
        &Longtime, &Hdr.h_namesize,&Longfile,Hdr.h_name);
```

Longtime and *Longfile* are equivalent to *Hdr.h_mtime* and *Hdr.h_filesize*, respectively. The contents of each file are recorded in an element of the array of varying length structures, *archive*, together with other items describing the file. Every instance of *h_magic* contains the constant 070707 (octal). The items *h_dev* through *h_mtime* have meanings explained in *stat*(2). The length of the null-terminated path name *h_name*, including the null byte, is given by *h_namesize*.

The last record of the *archive* always contains the name TRAILER!!!. Special files, directories, and the trailer are recorded with *h_filesize* equal to zero.

SEE ALSO

cpio(1), *find*(1), *stat*(2).

NAME

dir — format of directories

SYNOPSIS

```
#include <sys/dir.h>
```

DESCRIPTION

A directory behaves exactly like an ordinary file, save 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).

NAME

environ – user environment

SYNOPSIS

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

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*, *time*, *nice(1)*, etc., apply in searching for a file known by an incomplete path name. The prefixes are separated by `:`.

Login(1) sets :

```
PATH=:/bin:/usr/bin.
```

HOME A user's login directory, set by *login(1)* from the password file *passwd(5)*.

TERM The kind of terminal for which output is to be prepared. This information is used by commands, such as *nroff*, *more*, or *vi*, which may exploit special terminal capabilities. See */etc/termcap* or *(termcap(5))* for a list of terminal types.

SHELL The file name of the users login shell.

TERMCAP The string describing the terminal in **TERM**, or the name of the termcap file, see *termcap(5)*.

EXINIT A startup list of commands read by *ex(1)*, *edit(1)*, and *vi(1)*.

USER The login name of the user.

Further names may be placed in the environment by the *export* command and '*name=value*' arguments in *sh(1)*, or by the *setenv* command if you use *csh(1)*. Arguments may also be placed in the environment at the point of an *exec(2)*. It is unwise to conflict with certain *sh(1)* variables that are frequently exported by ".profile" files: MAIL, PS1, PS2, IFS.

SEE ALSO

csh(1), *ex(1)*, *login(1)*, *sh(1)*, *exec(2)*, *system(3S)*, *termcap(5)*, *tty(7)*.

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 (with header) */
    time_t     e_time;      /* time of day */
};
```

The permissible record types are as follows:

```
#define E_GOTS  010 /* Start for UNIX/TS */
#define E_GORT  011 /* Start for UNIX/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 */
```

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 {
    struct utsname e_name; /* system names */
    unsigned      e_bconf; /* block device configuration */
};

#define eend errhdr
struct etimchg {
    time_t      e_ntime; /* new time */
};
```

Stray interrupts cause a record with the following format to be logged in the file:

```
struct estray {
    physadr      e_saddr; /* stray loc or device addr */
    unsigned     e_sbacty; /* active block devices */
};
```

Memory parity error record that is logged whenever one occurs, hardware permitting:

```
struct eparity {
    int          e_parreg; /* memory subsystem registers */
};
```

Error records for block devices have the following format:

```
struct eblock {
    dev_t      e_dev;      /* "true" major + minor dev number */
    unsigned   e_bacty;    /* other block I/O activity */
    struct iostat e_stats; /* unit I/O statistics */
    short      e_bflags;   /* read/write, error, etc */
    short      e_nreg;     /* number of device registers */
    daddr_t    e_bnum;     /* logical block number */
    unsigned   e_bytes;    /* number of bytes to transfer */
    paddr_t    e_memadd;   /* buffer memory address */
    ushort     e_rtry;     /* number of retries where */
    struct pos {           /* the block device the error occurred */
        /* set invalid fields to -1 */
        unsigned unit;
        unsigned cyl;
        unsigned trk;
        unsigned sector;
    } e_pos;
};
```

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 */
```

SEE ALSO

errdemon(1M).

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_iseize;      /* 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 inodes */
    char       s_fname[6];   /* file system name */
    char       s_fpack[6];   /* file system pack name */
};

#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.

S_iseize 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_iseize* - 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 inodes (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 inode is really free or not is maintained in the inode itself.

S_tinode is the total free inodes 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_fmod* 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 inode and its flags, see *inode*(4).

FILES

/usr/include/sys/filsys.h
/usr/include/sys/stat.h

SEE ALSO

fsck(1M), fsdb(1M), mkfs(1M), inode(4).

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:

t*tabs* 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.

s*size* 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.

m*margin* 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* which may be used to convert files to a standard format acceptable to other UNIX System commands.

SEE ALSO

ed(1), newform(1), tabs(1).

NAME

gettydefs — speed and terminal settings used by getty

DESCRIPTION

The `/etc/gettydefs` file contains information used by `getty(1M)` (see the *UniPlus⁺ Administrator's Manual*) 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. Lines that begin with `#` are ignored and may be used to comment the file. 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:

<i>label</i>	This is the string against which <code>getty</code> tries to match its second argument. It is often the speed, such as <code>1200</code> , at which the terminal is supposed to run, but it needn't be (see below).
<i>initial-flags</i>	These flags are the initial <code>ioctl(2)</code> settings to which the terminal is to be set if a terminal type is not specified to <code>getty</code> . <code>Getty</code> understands the symbolic names specified in <code>/usr/include/sys/termio.h</code> (see <code>termio(7)</code> in the <i>UniPlus⁺ Administrator's Manual</i>). Normally only the speed flag is required in the <i>initial-flags</i> . <code>Getty</code> automatically sets the terminal to raw input mode and takes care of most of the other flags. The <i>initial-flag</i> settings remain in effect until <code>getty</code> executes <code>login(1)</code> .
<i>final-flags</i>	These flags take the same values as the <i>initial-flags</i> and are set just prior to <code>getty</code> executes <code>login</code> . The speed flag is again required. The composite flag <code>SANE</code> 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 <i>final-flags</i> are <code>TAB3</code> , so that tabs are sent to the terminal as spaces, and <code>HUPCL</code> , so that the line is hung up on the final close.
<i>login-prompt</i>	This entire field is printed as the <i>login-prompt</i> . Unlike the above fields where white space is ignored (a space, tab or new-line), they are included in the <i>login-prompt</i> field.
<i>next-label</i>	This indicates the next <i>label</i> of the entry in the table that <code>getty</code> should use if the user types a <code>< break ></code> or the input cannot be read. Usually, a series of speeds are linked together in this fashion, into a closed set. For instance, <code>2400</code> linked to <code>1200</code> , which in turn is linked to <code>300</code> , which finally is linked to <code>2400</code> .

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` can't 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.

GETTYDEFS (4)

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

`getty(1M)`, `termio(7)` in the *UniPlus+ Administrator's Manual*
`login(1)`, `ioctl(2)`.

GETTYDEFS (4)

GPS (4)

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 `star(1G)`) and `vloc` (in `loc(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:

lines *cw points sw*
arc *cw points sw*
text *cw point sw so [string]*
hardware *cw point [string]*
comment *cw [string]*

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. The first eight bits contain *color* information. In *arc* and *lines* the last eight bits are divided as four bits *weight* and four bits *style*. In the *text* primitive the last eight bits of *sw* contain the *font*.

so *So* is the size-orientation word used in *text* primitives. The first eight bits contain text size and the remaining 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.

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 user 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

newgrp(1), *passwd(1)*, *crypt(3C)*, *passwd(4)*.

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 newline, however, a backslash (\) preceding a newline 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 *getty*s 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 to four 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 *action* will be taken on this *process* for 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 the system is only in these states for as long as it takes to execute all the entries associated with the states. 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 and 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* is 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 can use **s** to specify that *init* start in the

SINGLE USER state. Additionally, if *init* doesn't 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

getty(1M), *init(1M)* in the *UniPlus+ Administrator's Manual*.
sh(1), *who(1)*, *exec(2)*, *open(2)*, *signal(2)*.

NAME

inode — format of an inode

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 created */
};
/*
 * 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)`.

NAME

master — master device information table

DESCRIPTION

This file is used by *config* to obtain device information that enables it to generate the configuration files. The file consists of 3 parts, each separated by a line with a dollar sign (\$) in column 1. Part 1 contains device information; part 2 contains names of devices that have aliases; part 3 contains tunable parameter information. Any line with an asterisk (*) in column 1 is treated as a comment.

Part 1 contains lines consisting of at least 10 fields and at most 13 fields, with the fields delimited by tabs and/or blanks:

Field 1: device name (8 chars. maximum).
 Field 2: interrupt vector size (decimal, in bytes).
 Field 3: device mask (octal) — each "on" bit indicates that the handler exists:
 000100 initialization handler
 000040 power-failure handler
 000020 open handler
 000010 close handler
 000004 read handler
 000002 write handler
 000001 ioctl handler.
 Field 4: device type indicator (octal):
 000200 allow only one of these devices
 000100 suppress count field in the *conf.c* file
 000040 suppress interrupt vector
 000020 required device
 000010 block device
 000004 character device
 000002 floating vector
 000001 fixed vector.
 Field 5: handler prefix (4 chars. maximum).
 Field 6: device address size (decimal).
 Field 7: major device number for block-type device.
 Field 8: major device number for character-type device.
 Field 9: maximum number of devices per controller (decimal).
 Field 10: maximum bus request level (4 through 7).
 Fields 11-13: optional configuration table structure declarations (8 chars. maximum).

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)

Devices that are not interrupt-driven have an interrupt vector size of zero. The 040 bit in Field 4 causes *config* to record the interrupt vector although the *ivec.s* file will show no interrupt vector assignment at those locations (interrupts here will be treated as strays).

NAME

mnttab — mounted file system table

SYNOPSIS

```
#include <mnttab.h>
```

DESCRIPTION

Mnttab 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[10];
        char      mt_filsys[10];
        short     mt_ro_flg;
        time_t    mt_time;
};
```

Each entry is 26 bytes in length; the first 10 bytes are the null-padded name of the place where the *special file* is mounted; the next 10 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.

FILES

/etc/mnttab

SEE ALSO

mount(1M), *setmnt(1M)*.

NAME

passwd — password file

DESCRIPTION*Passwd* contains for each user the following information:

login name
 encrypted password
 numerical user ID
 numerical group ID
 user's real name, and other information if desired
 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 ID's 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 his 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. A user who attempts to login after his password has 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 *..*) the user will be forced to change his password the next time he logs in (and the "age" will disappear from his entry in the password file). If *m* > *M* (signified, e.g., by the string *./*) only the super-user will be able to change the password.

FILES*/etc/passwd***SEE ALSO**

login(1), passwd(1), a64l(3C), crypt(3C), getpwent(3C), group(4).

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", "shortdashed", and "dotdashed". Effective only for the **-T4014** and **-Tver** 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

tplot(1G), *plot(3X)*, *gps(4)*, *term(5)*.

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.

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 the shell's **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

env(1), login(1), mail(1), sh(1), stty(1), su(1), environ(5), term(5).

NAME

sccsfile — 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:
@hDDDDDD

The value of the checksum is the sum of all characters, except those of the first line. The @h provides a *magic number* of (octal) 064001.

Delta table

The delta table consists of a variable number of entries of the form:

```
@s DDDDD/DDDD/DDDD
@d <type> <SCCS ID> r/mo/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.

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
@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. z flag is used in certain specialized interface programs.

Comments

Arbitrary text 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 don't 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)*.
Source Code Control System User's Guide

NAME

tp - magnetic tape format

DESCRIPTION

The command *tp(1)* dumps files to and extracts files from magtape.

Block zero contains a copy of a stand-alone bootstrap program.

Blocks 1 through 62 contain a directory of the tape. There are 496 entries in the directory; 8 entries per block; 64 bytes per entry. Each entry has the following format:

```

struct  tpent {
        char   pathnam[32];
        short  mode;
        char   uid;
        char   uid;
        char   gid;
        char   spare;
        char   size0;
        short  size2;
        long   time;
        short  tapea;          /* tape address */
        short  unused[8];
        short  cksum;         /* check sum */
}

```

The *pathnam* entry is the path name of the file when put on the tape. If the path name starts with a zero word, the entry is empty. It is at most 32 bytes long and ends in a null byte. *Mode*, *uid*, *gid*, the sizes and time modified are the same as described under i-nodes (*fs(4)*). The tape address is the tape block number of the start of the contents of the file. Every file starts on a block boundary. The file occupies $(\text{size}+511)/512$ blocks of continuous tape. The checksum entry has a value such that the sum of the 32 words of the directory entry is zero.

Blocks 63 on are available for file storage.

A fake entry has a size of zero. See *tp(1)*.

SEE ALSO

cpio(1), *tp(1)*, *fs(4)*.

NAME

ttytype — data base of terminal types by port

DESCRIPTION

Ttytype is a database containing, for each tty port on the system, the kind of terminal that is attached to it. There is one line per port, containing the terminal kind (as a name listed in *termcap(5)*), a space, and the name of the tty, minus */dev/*.

This information is read by *tset(1)* and by *login(1)* to initialize the TERM environment variable at login time.

EXAMPLE

```
dw console
3a tty0
h19 tty1
h19 tty2
du ttyd0
```

FILES

/etc/ttytype

SEE ALSO

tset(1), *login(1)*.

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, 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 */
};

/* 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
```

UTMP(4)

UTMP(4)

INTRO(5)

INTRO(5)

SEE ALSO

login(1), who(1), write(1), getut(3C).

NAME

intro — introduction to miscellany

DESCRIPTION

This section describes miscellaneous facilities such as macro packages, character set tables, etc.

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), by *setenv* in *csh*(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

env(1), *login*(1), *sh*(1), *exec*(2), *getenv*(3C), *profile*(4), *term*(5).

NAME

eqnchar — special character definitions for eqn and neqn

SYNOPSIS

eqn /usr/pub/eqnchar [files] | troff [options]

neqn /usr/pub/eqnchar [files] | nroff [options]

DESCRIPTION

Eqnchar contains *troff* and *nroff* character definitions for constructing characters that are not available on the Wang Laboratories, Inc. C/A/T phototypesetter. These definitions are primarily intended for use with *eqn* and *neqn*; *eqnchar* contains definitions for the following characters:

<i>ciplus</i>	\oplus			<i>square</i>	\square
<i>citimes</i>	\otimes	<i>langle</i>	\langle	<i>circle</i>	\circ
<i>wig</i>	\sim	<i>rangle</i>	\rangle	<i>blot</i>	\blacksquare
<i>-wig</i>	\cong	<i>hbar</i>	\hbar	<i>bullet</i>	\bullet
<i>>wig</i>	\succ	<i>ppd</i>	\perp	<i>prop</i>	\propto
<i><wig</i>	\preceq	<i><-></i>	\leftrightarrow	<i>empty</i>	\emptyset
<i>=wig</i>	\equiv	<i><=></i>	\Leftrightarrow	<i>member</i>	\in
<i>star</i>	*	<i> <</i>	\langle	<i>nomem</i>	\notin
<i>bigstar</i>	*	<i> ></i>	\rangle	<i>cup</i>	\cup
<i>=dot</i>	\doteq	<i>ang</i>	\sphericalangle	<i>cap</i>	\cap
<i>orsign</i>	\vee	<i>rang</i>	\sphericalangle	<i>incl</i>	\sqsubseteq
<i>andsign</i>	\wedge	<i>3dot</i>	\vdots	<i>subset</i>	\subset
<i>=del</i>	\triangleq	<i>thf</i>	\therefore	<i>supset</i>	\supset
<i>oppA</i>	\forall	<i>quarter</i>	$\frac{1}{4}$	<i>!subset</i>	$\not\subset$
<i>oppE</i>	\exists	<i>3quarter</i>	$\frac{3}{4}$	<i>!supset</i>	$\not\supset$
<i>angstrom</i>	\AA	<i>degree</i>	$^\circ$	<i>scrL</i>	\emptyset
<i>==<</i>	\leq	<i>==></i>	\geq		

FILES

/usr/pub/eqnchar

SEE ALSO

eqn(1), nroff(1), troff(1).

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 /* Non-blocking I/O */

#define O_APPEND 010 /* append (writes guaranteed at the end) */

/* 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 */

SEE ALSO

fcntl(2), open(2).

NAME

greek — graphics for the extended TTY-37 type-box

SYNOPSIS

cat /usr/pub/greek [| greek -Tterminal]

DESCRIPTION

Greek gives the mapping from ASCII to the “shift-out” graphics in effect between SO and SI on TELETYPE® Model 37 terminals equipped with a 128-character type-box. These are the default greek characters produced by *nroff*. The filters of *greek*(1) attempt to print them on various other terminals. The file contains:

alpha	α	A	beta	β	B	gamma	γ	\
GAMMA	Γ	G	delta	δ	D	DELTA	Δ	W
epsilon	ϵ	S	zeta	ζ	Q	eta	η	N
THETA	Θ	T	theta	θ	O	lambda	λ	L
LAMBDA	Λ	E	mu	μ	M	nu	ν	@
xi	ξ	X	pi	π	J	PI	Π	P
rho	ρ	K	sigma	σ	Y	SIGMA	Σ	R
tau	τ	I	phi	ϕ	U	PHI	Φ	F
psi	ψ	V	PSI	Ψ	H	omega	ω	C
OMEGA	Ω	Z	nabla	∇		not	\neg	-
partial	∂]	integral	\int	^			

FILES

/usr/pub/greek

SEE ALSO

300(1), 4014(1), 450(1), greek(1), tc(1), nroff(1).

NAME

inet — Internet protocol family

SYNOPSIS

```
#include <net/in.h>
```

DESCRIPTION

The Internet protocol family is a collection of protocols layered atop the *Internet Protocol (IP)* transport layer, and utilizing the Internet address format. The Internet family provides protocol support for the SOCK_STREAM, SOCK_DGRAM, and SOCK_RAW socket types; the SOCK_RAW interface provides access to the *IP* protocol.

ADDRESSING

Internet addresses are four byte quantities, stored in network standard format. The include file <net/in.h> defines this address as a discriminated union with the following conventions,

```
/*
 * Internet address
 */
struct in_addr {
    union {
        struct { u_char s_b1,s_b2,s_b3,s_b4; } S_un_b;
        struct { u_short s_w1,s_w2; } S_un_w;
        u_long S_addr;
    } S_un;
#define s_addr S_un.S_addr /* can be used for most tcp & ip code */
#define s_host S_un.S_un_b.s_b2 /* host on imp */
#define s_net S_un.S_un_b.s_b1 /* network */
#define s_imp S_un.S_un_w.s_w2 /* imp */
#define s_impno S_un.S_un_b.s_b4 /* imp # */
#define s_lh S_un.S_un_b.s_b3 /* logical host */
};
```

Sockets bound to the Internet protocol family utilize the following addressing structure,

```
struct sockaddr_in {
    short sin_family;
    u_short sin_port;
    struct in_addr sin_addr;
    char sin_zero[8];
};
```

Sockets may be created with the address INADDR_ANY to effect "wildcard" matching on incoming messages.

PROTOCOLS

The Internet protocol family is comprised of the *IP* transport protocol, Internet Control Message Protocol (*ICMP*), Transmission Control Protocol (*TCP*), and User Datagram Protocol (*UDP*). *TCP* is used to support the SOCK_STREAM abstraction while *UDP* is used to support the SOCK_DGRAM abstraction. A raw interface to *IP* is available by creating an Internet socket of type SOCK_RAW. The *ICMP* message protocol is not directly accessible.

INTERFACES

A number of interfaces are usable with the Internet protocol family. These

include various Ethernet interfaces and standard a "software loopback" interface.

SEE ALSO

ip(5), lo(5) tcp(5), udp(5).

NAME

ip — Internet Protocol

SYNOPSIS

```
#include <net/socket.h>
#include <net/in.h>
```

```
struct sockproto proto = { PF_INET, ? };
```

```
socket(SOCK_RAW, &proto, address, options);
struct sockaddr_in *address; int options;
```

DESCRIPTION

IP is the transport layer protocol used by the Internet protocol family. It may be accessed through a "raw socket" when developing new protocols, or special purpose applications. *IP* sockets are connectionless, and are normally used with the *send(2)* and *receive(2)* calls, though the *connect(2)* call may also be used to fix the destination for future packets (in which case the *read(2)* and *write(2)* system calls may be used).

Outgoing packets automatically have an *IP* header prepended to them (based on the destination address and the protocol number the socket is created with). Likewise, incoming packets have their *IP* header stripped before being sent to the user. It is currently not possible to send or receive *IP* options.

DIAGNOSTICS

EISCONN when trying to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected;

ENOTCONN when trying to send a datagram, but no destination address is specified, and the socket hasn't been connected;

ENOBUFS when the system runs out of memory for an internal data structure;

EADDRNOTAVAIL when an attempt is made to create a socket with a network address for which no network interface exists.

SEE ALSO

inet(5), net(5).

BUGS

One should be able to send and receive ip options.

The protocol should be settable after socket creation.

NAME

loop — software loopback interface

SYNOPSIS

pseudo-device loop

DESCRIPTION

The *loop* interface is a software loopback mechanism which may be used for performance analysis, software testing, and/or local communication. The interface is Internet addressable as network 127 (decimal). The local host is host 1.

DIAGNOSTICS

lo%d: can't handle af%d

The interface was handed a message with addresses formatted in an unsuitable address family; the packet was dropped.

SEE ALSO

inet(5), net(5).

BUGS

It should handle all address and protocol families.

NAME

man — macros for formatting entries in this manual

SYNOPSIS

nroff — man files

troff — man [-rs1] files

DESCRIPTION

These *troff*(1) macros are used to lay out the format of the entries of this manual. These macros are used by the *man*(1) command.

The default page size is 8.5"×11", with a 6.5"×10" text area; the *-rs1* option reduces these dimensions to 6"×9" and 4.75"×8.375", respectively; this option (which is *not* effective in *nroff*) also reduces the default type size from 10-point to 9-point, and the vertical line spacing from 12-point to 10-point. The *-rV2* option may be used to set certain parameters to values appropriate for certain Versatec printers: it sets the line length to 82 characters, the page length to 84 lines, and it inhibits underlining; this option should not be confused with the *-Tvp* option of the *man*(1) command, which is available at some UNIX System sites.

Any *text* argument below may be one to six "words". Double quotes (") may be used to include blanks in a "word". If *text* is empty, the special treatment is applied to the next line that contains text to be printed. For example, *.I* may be used to italicize a whole line, or *.SM* followed by *.B* to make small bold text. By default, hyphenation is turned off for *nroff*, but remains on for *troff*.

Type font and size are reset to default values before each paragraph and after processing font- and size-setting macros, e.g., *.I*, *.RB*, *.SM*. Tab stops are neither used nor set by any macro except *.DT* and *.TH*.

Default units for indents *in* are ens. When *in* is omitted, the previous indent is used. This remembered indent is set to its default value (7.2 ens in *troff*, 5 ens in *nroff*—this corresponds to 0.5" in the default page size) by *.TH*, *.P*, and *.RS*, and restored by *.RE*.

.TH t s c n Set the title and entry heading; *t* is the title, *s* is the section number, *c* is extra commentary, e.g., "local", *n* is new manual name. Invokes *.DT* (see below).

.SH text Place subhead *text*, e.g., SYNOPSIS, here.

.SS text Place sub-subhead *text*, e.g., Options, here.

.B text Make *text* bold.

.I text Make *text* italic.

.SM text Make *text* 1 point smaller than default point size.

.RI a b Concatenate roman *a* with italic *b*, and alternate these two fonts for up to six arguments. Similar macros alternate between any two of roman, italic, and bold:

.IR .RB .BR .IB .BI

.P Begin a paragraph with normal font, point size, and indent. *.PP* is a synonym for *.P*.

.HP in Begin paragraph with hanging indent.

.TP in Begin indented paragraph with hanging tag. The next line that contains text to be printed is taken as the tag. If the tag does not fit, it is printed on a separate line.

.IP t in Same as *.TP in* with tag *t*; often used to get an indented paragraph without a tag.

- .RS** *in* Increase relative indent (initially zero). Indent all output an extra *in* units from the current left margin.
- .RE** *k* Return to the *k*th relative indent level (initially, *k*=1; *k*=0 is equivalent to *k*=1); if *k* is omitted, return to the most recent lower indent level.
- .PM** *m* Produces proprietary markings; where *m* may be **P** for **PRIVATE**, **N** for **NOTICE**, **BP** for **BELL LABORATORIES PROPRIETARY**, or **BR** for **BELL LABORATORIES RESTRICTED**.
- .DT** Restore default tab settings (every 7.2 ens in *troff*, 5 ens in *nroff*).
- .PD** *v* Set the interparagraph distance to *v* vertical spaces. If *v* is omitted, set the interparagraph distance to the default value (0.4*v* in *troff*, 1*v* in *nroff*).

The following *strings* are defined:

- *R** ® in *troff*, (**Reg.**) in *nroff*.
- *S** Change to default type size.
- *(Tm** Trademark indicator.

The following *number registers* are given default values by **.TH**:

- IN** Left margin indent relative to subheads (default is 7.2 ens in *troff*, 5 ens in *nroff*).
- LL** Line length including **IN**.
- PD** Current interparagraph distance.

CAVEATS

In addition to the macros, strings, and number registers mentioned above, there are defined a number of *internal* macros, strings, and number registers. Except for names predefined by *troff* and number registers **d**, **m**, and **y**, all such internal names are of the form *XA*, where *X* is one of **(**, **l**, and **]**, and *A* stands for any alphanumeric character.

If a manual entry needs to be preprocessed by *cw*(1), *eqn*(1) (or *neqn*), and/or *tbl*(1), it must begin with a special line (described in *man*(1)), causing the *man* command to invoke the appropriate preprocessor(s).

The programs that prepare the Table of Contents and the Permuted Index for this Manual assume the *NAME* section of each entry consists of a single line of input that has the following format:

```
name[, name, name ...] \- explanatory text
```

The macro package increases the inter-word spaces (to eliminate ambiguity) in the *SYNOPSIS* section of each entry.

The macro package itself uses only the roman font (so that one can replace, for example, the bold font by the constant-width font—see *cw*(1)). Of course, if the input text of an entry contains requests for other fonts (e.g., **.I**, **.RB**, **\fI**), the corresponding fonts must be mounted.

EXAMPLE

```
nroff -man man.5
```

to *nroff* this manual section.

FILES

```
/usr/lib/tmac/tmac.an
/usr/lib/macros/cmp.[nt].[dt].an
```

```
/usr/lib/macros/ucmp.[nt].an
/usr/man/[ua]_man/man0/skeleton
```

SEE ALSO

man(1), *nroff*(1), *troff*(1).

BUGS

If the argument to **.TH** contains *any* blanks and is *not* enclosed by double quotes (**), there will be bird-dropping-like things on the output.

NAME

mm — the MM macro package for formatting documents

SYNOPSIS

mm [options] [files]

nroff -mm [options] [files]

nroff -cm [options] [files]

mmt [options] [files]

troff -mm [options] [files]

troff -cm [options] [files]

DESCRIPTION

This package provides a formatting capability for a very wide variety of documents. It is the standard package used by the BTL typing pools and documentation centers. The manner in which a document is typed in and edited is essentially independent of whether the document is to be eventually formatted at a terminal or is to be phototypeset. See the references below for further details.

The `-mm` option causes `nroff` and `troff(1)` to use the non-compacted version of the macro package, while the `-cm` option results in the use of the compacted version, thus speeding up the process of loading the macro package.

FILES

/usr/lib/tmac/tmac.m	pointer to the non-compacted version of the package
/usr/lib/macros/mm[nt]	non-compacted version of the package
/usr/lib/macros/cmp.[nt].[dt].m	compacted version of the package
/usr/lib/macros/ucmp.[nt].m	initializers for the compacted version of the package

SEE ALSO

mm(1), mmt(1), nroff(1), troff(1).

MM—Memorandum Macros by D. W. Smith and J. R. Mashey.

Typing Documents with MM by D. W. Smith and E. M. Piskorik.

NAME

mosd — the OSDD adapter macro package for formatting documents

SYNOPSIS

```
osdd [ options ] [ files ]
mm -mosd [ options ] [ files ]
nroff -mm -mosd [ options ] [ files ]
nroff -cm -mosd [ options ] [ files ]

mmt -mosd [ options ] [ files ]
troff -mm -mosd [ options ] [ files ]
troff -cm -mosd [ options ] [ files ]
```

DESCRIPTION

The OSDD adapter macro package is a tool used in conjunction with the MM macro package to prepare Operations Systems Deliverable Documentation. Many of the OSDD Standards are different than the default format provided by MM. The OSDD adapter package sets the appropriate MM options for automatic production of the OSDD Standards. The OSDD adapter package also generates the correct OSDD page headers and footers, heading styles, Table of Contents format, etc.

OSDD document (input) files are prepared with the MM macros. Additional information which must be given at the beginning of the document file is specified by the following string definitions:

```
.ds H1 document-number
.ds H2 section-number
.ds H3 issue-number
.ds H4 date
.ds H5 rating
```

The *document-number* should be of the standard 10 character format. The words "Section" and "Issue" should not be included in the string definitions; they will be supplied automatically when the document is printed. For example:

```
.ds H1 OPA-1P135-01
.ds H2 4
.ds H3 2
```

automatically produces

```
OPA-1P135-01
Section 4
Issue 2
```

as the document page header. Quotation marks are not used in string definitions.

If certain information is not to be included in a page header, then the string is defined as null; e.g.,

```
.ds H2
```

means that there is no *section-number*.

The OSDD Standards require that the *Table of Contents* be numbered beginning with *Page 1*. By default, the first page of text will be numbered *Page 2*. If the *Table of Contents* has more than one page, for example *n*, then either `-rPn+1` must be included as a command line option or `.nr P n` must be included in the document file. For example, if the *Table of*

Contents is four pages then use `-rP5` on the command line or `.nr P 4` in the document file.

The OSDD Standards require that certain information such as the document *rating* appear on the *Document Index* or on the *Table of Contents* page if there is no index. By default, it is assumed that an index has been prepared separately. If there is no index, the following must be included in the document file:

```
.nr Di 0
```

This will ensure that the necessary information is included on the *Table of Contents* page.

The OSDD Standards require that all numbered figures be placed at the end of the document. The `.Fg` macro is used to produce full page figures. This macro produces a blank page with the appropriate header, footer, and figure caption. Insertion of the actual figure on the page is a manual operation. The macro usage is

```
.Fg page-count "figure caption"
```

where *page-count* is the number of pages required for a multi-page figure (default 1 page).

Figure captions are produced by the `.Fg` macro using the `.BS/.BE` macros. Thus the `.BS/.BE` macros are also not available for users. The `.Fg` macro cannot be used within the document unless the final `.Fg` in a series of figures is followed by a `.SK` macro to force out the last figure page.

The *Table of Contents* for OSDD documents (see Figure 4 in Section 4.1 of the OSDD Standards) is produced with:

```
.Tc
System Type
System Name
Document Type
.Td
```

The `.Tc/.Td` macros are used instead of the `.TC` macro from MM.

By default, the adapter package causes the NOTICE disclosure statement to be printed. The `.PM` macro may be used to suppress the NOTICE or to replace it with the PRIVATE disclosure statement as follows:

```
.PM      none printed
.PM P    PRIVATE printed
.PM N    NOTICE printed (default)
```

The `.P` macro is used for paragraphs. The `Np` register is set automatically to indicate the paragraph numbering style. It is very important that the `.P` macro be used correctly. All paragraphs (including those immediately following a `.H` macro) must use a `.P` macro. Unless there is a `.P` macro, there will not be a number generated for the paragraph. Similarly, the `.P` macro should not be used for text which is not a paragraph. The `.SP` macro may be appropriate for these cases, e.g., for "paragraphs" within a list item.

The page header format is produced automatically in accordance with the OSDD Standards. The OSDD Adapter macro package uses the `.TP` macro for this purpose. Therefore the `.TP` macro normally available in MM is not available for users.

FILES

```
/usr/lib/tmac/tmac.osd
```

SEE ALSO

mm(1), mmt(1), nroff(1), troff(1), mm(5).

MM—Memorandum Macros by D. W. Smith and J. R. Mashey.

Operations Systems Deliverable Documentation Standards, June 1980.

NAME

mptx — the macro package for formatting a permuted index

SYNOPSIS

nroff -mptx [options] [files]

troff -mptx [options] [files]

DESCRIPTION

This package provides a definition for the .xx macro used for formatting a permuted index as produced by *ptx(1)*. This package does not provide any other formatting capabilities such as headers and footers. If these or other capabilities are required, the *mptx* macro package may be used in conjunction with the *MM* macro package. In this case, the **-mptx** option must be invoked *after* the **-mm** call. For example:

nroff -cm -mptx file

or

mm -mptx file

FILES

/usr/lib/tmac/tmac.ptx pointer to the non-compacted version of the package

/usr/lib/macros/ptx non-compacted version of the package

SEE ALSO

mm(1), nroff(1), ptx(1), troff(1), mm(5).

NAME

mv — a troff macro package for typesetting view graphs and slides

SYNOPSIS

mvt [-a] [options] [files]

troff [-a] [-rX1] -mv [options] [files]

DESCRIPTION

This package makes it easy to typeset view graphs and projection slides in a variety of sizes. A few macros (briefly described below) accomplish most of the formatting tasks needed in making transparencies. All of the facilities of *troff*(1), *cw*(1), *eqn*(1), and *tbl*(1) are available for more difficult tasks.

The output can be previewed on most terminals, and, in particular, on the Tektronix 4014, as well as on the Versatec printer. For these two devices, specify the *-rX1* option (this option is automatically specified by the *mvt* command—q.v.—when that command is invoked with the *-T4014* or *-Tvp* options). To preview output on other terminals, specify the *-a* option.

The available macros are:

<p>.VS [n] [i] [d]</p>	<p>Foil-start macro; foil size is to be 7"×7"; <i>n</i> is the foil number, <i>i</i> is the foil identification, <i>d</i> is the date; the foil-start macro resets all parameters (indent, point size, etc.) to initial default values, except for the values of <i>i</i> and <i>d</i> arguments inherited from a previous foil-start macro; it also invokes the .A macro (see below).</p>
<p>The naming convention for this and the following eight macros is that the first character of the name (V or S) distinguishes between view graphs and slides, respectively, while the second character indicates whether the foil is square (S), small wide (w), small high (h), big wide (W), or big high (H). Slides are "skinnier" than the corresponding view graphs: the ratio of the longer dimension to the shorter one is larger for slides than for view graphs. As a result, slide foils can be used for view graphs, but not vice versa; on the other hand, view graphs can accommodate a bit more text.</p>	
<p>.Vw [n] [i] [d]</p>	<p>Same as .VS, except that foil size is 7" wide × 5" high.</p>
<p>.Vh [n] [i] [d]</p>	<p>Same as .VS, except that foil size is 5"×7".</p>
<p>.VW [n] [i] [d]</p>	<p>Same as .VS, except that foil size is 7"×5.4".</p>
<p>.VH [n] [i] [d]</p>	<p>Same as .VS, except that foil size is 7"×9".</p>
<p>.Sw [n] [i] [d]</p>	<p>Same as .VS, except that foil size is 7"×5".</p>
<p>.Sh [n] [i] [d]</p>	<p>Same as .VS, except that foil size is 5"×7".</p>
<p>.SW [n] [i] [d]</p>	<p>Same as .VS, except that foil size is 7"×5.4".</p>
<p>.SH [n] [i] [d]</p>	<p>Same as .VS, except that foil size is 7"×9".</p>
<p>.A [x]</p>	<p>Place text that follows at the first indentation level (left margin); the presence of <i>x</i> suppresses the ½ line spacing from the preceding text.</p>
<p>.B [m [s]]</p>	<p>Place text that follows at the second indentation level; text is preceded by a mark; <i>m</i> is the mark (default is a large bullet); <i>s</i> is the increment or decrement to the point size of the mark with respect to the <i>prevailing</i></p>

- point size (default is 0); if *s* is 100, it causes the point size of the mark to be the same as that of the *default* mark.
- .C** [*m* [*s*]] Same as **.B**, but for the third indentation level; default mark is a dash.
- .D** [*m* [*s*]] Same as **.B**, but for the fourth indentation level; default mark is a small bullet.
- .T** *string* *String* is printed as an over-size, centered title.
- .I** [*in*] [*a* [*x*]] Change the current text indent (does not affect titles); *in* is the indent (in inches unless dimensioned, default is 0); if *in* is signed, it is an increment or decrement; the presence of *a* invokes the **.A** macro (see below) and passes *x* (if any) to it.
- .S** [*p*] [*l*] Set the point size and line length; *p* is the point size (default is "previous"); if *p* is 100, the point size reverts to the *initial* default for the current foil-start macro; if *p* is signed, it is an increment or decrement (default is 18 for **.VS**, **.VH**, and **.SH**, and 14 for the other foil-start macros); *l* is the line length (in inches unless dimensioned; default is 4.2" for **.Vh**, 3.8" for **.Sh**, 5" for **.SH**, and 6" for the other foil-start macros).
- .DF** *n f* [*n f*...] Define font positions; may not appear within a foil's input text (i.e., it may only appear after all the input text for a foil, but before the next foil-start macro); *n* is the position of font *f*; up to four "*n f*" pairs may be specified; the first font named becomes the *prevailing* font; the initial setting is (**H** is a synonym for **G**):
.DF 1 H 2 I 3 B 4 S
- .DV** [*a*] [*b*] [*c*] [*d*] Alter the vertical spacing between indentation levels; *a* is the spacing for **.A**, *b* is for **.B**, *c* is for **.C**, and *d* is for **.D**; all non-null arguments must be dimensioned; null arguments leave the corresponding spacing unaffected; initial setting is:
.DV .5v .5v .5v 0v
- .U** *str1* [*str2*] Underline *str1* and concatenate *str2* (if any) to it.

The last four macros in the above list do not cause a break; the **.I** macro causes a break only if it is invoked with more than one argument; all the other macros cause a break.

The macro package also recognizes the following upper-case synonyms for the corresponding lower-case *troff* requests:

.AD .BR .CE .FI .HY .NA .NF .NH .NX .SO .SP .TA .TI

The **Tm** string produces the trademark symbol.

The input tilde (~) character is translated into a blank on output.

See the user's manual cited below for further details.

FILES

/usr/lib/tmac/tmac.v
 /usr/lib/macros/vmca

SEE ALSO

cw(1), eqn(1), mmt(1), tbl(1), troff(1).
A Macro Package for View Graphs and Slides by T. A. Dolotta and

D. W. Smith.

BUGS

The **.VW** and **.SW** foils are meant to be 9" wide by 7" high, but because the typesetter paper is generally only 8" wide, they are printed 7" wide by 5.4" high and have to be enlarged by a factor of 9/7 before use as view graphs; this makes them less than totally useful.

NAME

net — introduction to networking facilities

SYNOPSIS

```
#include <net/socket.h>
#include <net/route.h>
```

DESCRIPTION

This section briefly describes the networking facilities available on the system.

All network protocols are associated with a specific *protocol-family*. A protocol-family provides basic services to the protocol implementation to allow it function within a specific network environment. These services may include packet fragmentation and reassembly, routing, addressing, and basic transport. A protocol-family may support multiple methods of addressing, though the current protocol implementations do not. A protocol-family is normally comprised of a number of protocols, one per *socket*(2) type. It is not required that a protocol-family support all socket types. A protocol-family may contain multiple protocols supporting the same socket abstraction.

A protocol supports one of the socket abstractions detailed in *socket*(2). A specific protocol may be accessed either by creating a socket of the appropriate type and protocol-family, or by requesting the protocol explicitly when creating a socket. Protocols normally accept only one type of address format, usually determined by the addressing structure inherent in the design of the protocol-family/network architecture. Certain semantics of the basic socket abstractions are protocol specific. All protocols are expected to support the basic model for their particular socket type, but may, in addition, provide non-standard facilities or extensions to a mechanism. For example, a protocol supporting the SOCK_STREAM abstraction may allow more than one byte of out-of-band data to be transmitted per out-of-band message.

A network interface is similar to a device interface. Network interfaces comprise the lowest layer of the networking subsystem, interacting with the actual transport hardware. An interface may support one or more protocol families, and/or address formats.

PROTOCOLS

The following protocol family identifiers are in use,

```
#define PF_UNSPEC 0 /* unspecified */
#define PF_INET 2 /* internetwork: UDP, TCP, etc. */
#define PF_PUP 4 /* pup protocols: e.g., BSP */
```

One must be specified in the *sockproto* structure supplied at socket creation time,

```
struct sockproto {
    short    sp_family; /* protocol family */
    short    sp_protocol; /* protocol within family */
};
```

ADDRESSING

The following address formats are in use:

```
#define AF_UNSPEC 0 /* unspecified */
#define AF_INET 2 /* internetwork: UDP, TCP, etc. */
```



```
#define AF_PUP      4 /* pup protocols: e.g. BSP */
```

ROUTING

The network facilities provided limited packet routing. A simple set of data structures comprise a "routing table" used in selecting the appropriate network interface when outputting packets. This table contains a single entry for each route to a specific network or host. A user process, the routing daemon, maintains this data base with the aid of three socket specific *ioctl* (2) commands: SIOCADDRT, SIOCDELRT, SIOCCHGRT. The commands allow the addition, deletion, or change of a single routing table entry, respectively. Routing table manipulations may only be carried out by super user and are subject to certain restrictions. The restrictions are:

1. No identical entries may be present.
2. No entry may be deleted or changed while the entry is in use (to be explained further below).

A routing table entry has the following form, as defined in `< net/route.h >`;

```
struct rtenry {
    u_long    rt_hash;
    struct    sockaddr rt_dst;
    struct    sockaddr rt_gateway;
    short     rt_flags;
    short     rt_refcnt;
    u_long    rt_use;
    struct    ifnet *rt_ifp;
};
```

with *rt_flags* defined from,

```
#define RTF_UP      0x1 /* route is available for use */
#define RTF_DIRECT 0x2 /* direct route to destination */
#define RTF_HOST    0x4 /* route to a host, not network */
```

Routing table entries come in two flavors, for a specific host or for all hosts on a specific network. When the system is booted, each network interface which configures itself installs a routing table entry when it wishes to have packets sent through it. Normally the interface specifies the route through it is a "direct" connection to the destination host or network. If the route is direct, the transport layer of a protocol family usually requests the packet be sent to the same host specified in the packet. Otherwise, the interface may be requested to address the packet to an entity different from the eventual recipient (i.e., the packet is forwarded).

Routing table entries installed by a user process may not specify the hash, reference count, use, or interface fields; these are filled in by the routing routines. In addition, a request to delete or change an existing routing table entry may be denied or partially performed depending on the state of the route. If a route is currently in use (the reference count field is non-zero), a request to delete the entry will result in the route being marked "down" and the error EBUSY returned. If the route was to be changed, but it was in use, only the flags value is updated and the error EBUSY is returned. These semantics are intended to allow a routing daemon to invalidate an entry, await freeing of the entry from use, then modify it at a later time.

The routing code may return EEXIST if requested to add an already existent entry, ESRCH if requested to delete or change an entry and it couldn't be

found, or ENOBUFS if requested to add an entry and the system was low on resources.

There currently is no support for reading the routing tables; user processes are expected to read the kernel's memory through `/dev/kmem`.

The use field is used by the routing code in providing a simple round-robin scheme of route selection when multiple routes to a destination are present; the heuristic is to choose the least used route.

SEE ALSO

socket(2).

NAME

regex - 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 <regex.h>

char *compile(instring, expbuf, endbuf, eof)
char *instring, *expbuf, *endbuf;

int step(string, expbuf)
char *string, *expbuf;
```

DESCRIPTION

This page describes general purpose regular expression matching routines in the form of *ed(1)*, defined in `/usr/include/regex.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 <regex.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 *c* 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(c)` is always ignored.

RETURN(pointer)

This macro is used on normal exit of the *compile* routine. The value of the argument *pointer* 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 *compile* routine. The argument *val* is an error number (see table below for meanings). This call should never return.

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 one, 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 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 only try to match the regular expression to the beginning of the string. 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 a one 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 routines *ecmp* and *getrange* are trivial and are called by the routines previously mentioned.

EXAMPLE

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>
...
...
compile(*argv, expbuf, &expbuf[ESIZE], '\0');
...
if(step(linebuf, expbuf))
    succeed();
```

FILES

/usr/include/regexp.h

SEE ALSO

ed(1), *grep(1)*, *sed(1)*.

BUGS

The handling of *circf* is kludgy.
 The routine *ecmp* is equivalent to the Standard I/O routine *strncmp* and should be replaced by that routine.
 The actual code is probably easier to understand than this manual page.

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_IWWRITE 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

tcp — Internet Transmission Control Protocol

SYNOPSIS

```
#include <net/socket.h>
#include <net/in.h>
```

```
struct sockproto proto = { PF_INET, IPPROTO_TCP };
socket(SOCK_STREAM, &proto, address, options);
struct sockaddr_in *address; int options;
```

DESCRIPTION

The *TCP* protocol provides reliable, flow-controlled, two-way transmission of data. It is a byte-stream protocol used to support the *SOCK_STREAM* abstraction. *TCP* uses the standard Internet address format and, in addition, provides a per-host collection of "port addresses". Thus, each address is composed of an Internet address specifying the host and network, with a specific *TCP* port on the host identifying the peer entity.

Sockets utilizing the *TCP* protocol are either "active" or "passive". Active sockets initiate connections to passive sockets. By default *TCP* sockets are created active; to create a passive socket the *SO_ACCEPTCONN* option must be supplied. Only passive sockets may use the *accept(2)* call to accept incoming connections. Only active sockets may use the *connect(2)* call to initiate connections.

Passive sockets may "underspecify" their location to match incoming connection requests from multiple networks. This technique, termed "wildcard addressing", allows a single server to provide service to clients on multiple networks. To create a socket which listens on all networks, the Internet address *INADDR_ANY* is specified. The *TCP* port may still be specified at this time; if the port is not specified, the system will assign one. Once a connection has been established the socket's address is fixed by the peer entity's location. The address assigned the socket is the address associated with the network interface through which packets are being transmitted and received. Normally this address corresponds to the peer entity's network.

OPTIONS

The *TCP* implementation supports two non-standard features: "keep-alives" and "true out-of-band" data transmission.

Keep-alives are a mechanism used to check if a peer entity is still functional. This is implemented by periodically "polling" the remote machine if the connection has been idle. The current implementation transmits keep-alive packets on a connection which has been idle for longer than 1 minute. If, despite the keep-alive packets, no response has been seen within 4 minutes, the connection is aborted. This mechanism applies only to connection in an "established" state; if a connection is idle for 1 minute but not yet established, it is simply aborted. The keep-alive mechanism is enabled by creating a socket with the *SO_KEEPALIVE* option. [N.B.: *TCP* implementations which do not closely follow the *TCP* specification may not respond to keep-alive messages, causing connections to be closed without reason; in this case keep-alives should not be used]

In order to transmit "true" out-of-band data, the *SO_TRUEOOB* option may be specified. This facility requires cooperation by the peer to function properly; this is negotiated through *TCP* options at the time a connection is

established. When this mechanism is used, one byte of data may be sent as an urgent, high-priority message to the peer. This data utilizes a separate, out-of-band data sequence space and is not subject to the normal flow control mechanisms imposed by *TCP*. In addition, the data stream is also marked to indicate the point at which the out-of-band data was sent. A process may send out-of-band data with the *SIOCSENDOOb* call,

```
ioctl(fd, SIOCSENDOOb, &data);
```

and receive out-of-band data with the *SIOCRCVOOb* call,

```
ioctl(fd, SIOCRCVOOb, &data);
```

To find out if the read pointer is at the mark in the data stream, the *SIOCATMARK* call may be used,

```
ioctl(fd, SIOCATMARK, &yesno);
```

The variable *yesno* will be a 1 if the read pointer currently points at the mark, and 0 otherwise.

DIAGNOSTICS

EISCONN when trying to establish a connection on a socket which already has one;

ENOBUFS when the system runs out of memory for an internal data structure;

ETIMEDOUT when a connection was dropped due to excessive retransmissions;

ECONNRESET when the remote peer forces the connection to be closed;

ECONNREFUSED when the remote peer actively refuses connection establishment (usually because no process is listening to the port);

EADDRINUSE when an attempt is made to create a socket with a port which has already been allocated;

EADDRNOTAVAIL when an attempt is made to create a socket with a network address for which no network interface exists.

SEE ALSO

inet(5), *net(5)*.

BUGS

Value added "features" such as "keep-alives" and "true" out-of-band are experimental and not part of the protocol standard.

NAME

term — conventional names for terminals

DESCRIPTION

These names are used by certain commands (e.g., *nroff*, *mm(1)*, *man(1)*, *tabs(1)*) and are maintained as part of the shell environment (see *sh(1)*, *profile(4)*, and *environ(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® Terminal Model 33 KSR
37	TELETYPE Terminal Model 37 KSR
40-2	TELETYPE Terminal Model 40/2
40-4	TELETYPE Terminal Model 40/4
4540	TELETYPE Terminal 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	General Electric TermiNet 1200
tn300	General 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*.

See */etc/termcap* on your system for a complete list.

SEE ALSO

mm(1), *nroff*(1), *sh*(1), *stty*(1), *tabs*(1), *tplot*(1G), *profile*(4), *environ*(5).

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

termcap — terminal capability data base

SYNOPSIS

/etc/termcap

DESCRIPTION

Termcap is a data base describing terminals used, e.g., by *vi*(1). Terminals are described in *termcap* by giving a set of capabilities which they have, and by describing how operations are performed. Padding requirements and initialization sequences are included in *termcap*.

Entries in *termcap* consist of a number of ':' separated fields. The first entry for each terminal gives the names which are known for the terminal, separated by '|' characters. The first name is always 2 characters long and is used by older version 6 systems which store the terminal type in a 16 bit word in a systemwide data base. The second name given is the most common abbreviation for the terminal, and the last name given should be a long name fully identifying the terminal. The second name should contain no blanks; the last name may well contain blanks for readability.

CAPABILITIES

(P) indicates padding may be specified

(P*) indicates that padding may be based on no. lines affected

Name Type Pad? Description

ae	str	(P)	End alternate character set
al	str	(P*)	Add new blank line
am	bool		Terminal has automatic margins
as	str	(P)	Start alternate character set
bc	str		Backspace if not ^H
bs	bool		Terminal can backspace with ^H
bt	str	(P)	Back tab
bw	bool		Backspace wraps from column 0 to last column
CC	str		Command character in prototype if terminal settable
cd	str	(P*)	Clear to end of display
ce	str	(P)	Clear to end of line
ch	str	(P)	Like cm but horizontal motion only, line stays same
cl	str	(P*)	Clear screen
cm	str	(P)	Cursor motion
co	num		Number of columns in a line
cr	str	(P*)	Carriage return, (default ^M)
cs	str	(P)	Change scrolling region (vt100), like cm
cv	str	(P)	Like ch but vertical only.
da	bool		Display may be retained above
dB	num		Number of millisecc of bs delay needed
db	bool		Display may be retained below
dC	num		Number of millisecc of cr delay needed
dc	str	(P*)	Delete character
dF	num		Number of millisecc of ff delay needed
dl	str	(P*)	Delete line
dm	str		Delete mode (enter)
dN	num		Number of millisecc of nl delay needed
do	str		Down one line
dT	num		Number of millisecc of tab delay needed
ed	str		End delete mode

ei	str	End insert mode; give :ei=: if ic
eo	str	Can erase overstrikes with a blank
ff	str	(P*) Hardcopy terminal page eject (default ^L)
hc	bool	Hardcopy terminal
hd	str	Half-line down (forward 1/2 linefeed)
ho	str	Home cursor (if no cm)
hu	str	Half-line up (reverse 1/2 linefeed)
hz	str	Hazeltine; can't print ~'s
ic	str	(P) Insert character
if	str	Name of file containing is
im	str	Insert mode (enter); give :im=: if ic
in	bool	Insert mode distinguishes nulls on display
ip	str	(P*) Insert pad after character inserted
is	str	Terminal initialization string
k0-k9	str	Sent by other function keys 0-9
kb	str	Sent by backspace key
kd	str	Sent by terminal down arrow key
ke	str	Out of keypad transmit mode
kh	str	Sent by home key
kl	str	Sent by terminal left arrow key
kn	num	Number of other keys
ko	str	Termcap entries for other non-function keys
kr	str	Sent by terminal right arrow key
ks	str	Put terminal in keypad transmit mode
ku	str	Sent by terminal up arrow key
l0-l9	str	Labels on other function keys
li	num	Number of lines on screen or page
ll	str	Last line, first column (if no cm)
ma	str	Arrow key map, used by vi version 2 only
mi	bool	Safe to move while in insert mode
ml	str	Memory lock on above cursor.
ms	bool	Safe to move while in standout and underline mode
mu	str	Memory unlock (turn off memory lock).
nc	bool	No correctly working carriage return (DM2500,H2000)
nd	str	Non-destructive space (cursor right)
nl	str	(P*) Newline character (default \n)
ns	bool	Terminal is a CRT but doesn't scroll.
os	bool	Terminal overstrikes
pc	str	Pad character (rather than null)
pt	bool	Has hardware tabs (may need to be set with is)
se	str	End stand out mode
sf	str	(P) Scroll forwards
sg	num	Number of blank chars left by so or se
so	str	Begin stand out mode
sr	str	(P) Scroll reverse (backwards)
ta	str	(P) Tab (other than ^I or with padding)
tc	str	Entry of similar terminal - must be last
te	str	String to end programs that use cm
ti	str	String to begin programs that use cm
uc	str	Underscore one char and move past it
ue	str	End underscore mode
ug	num	Number of blank chars left by us or ue
ul	bool	Terminal underlines even though it doesn't overstrike

up	str	Upline (cursor up)
us	str	Start underscore mode
vb	str	Visible bell (may not move cursor)
ve	str	Sequence to end open/visual mode
vs	str	Sequence to start open/visual mode
xb	bool	Beehive (f1=escape, f2=ctrl C)
xn	bool	A newline is ignored after a wrap (Concept)
xr	bool	Return acts like ce \r \n (Delta Data)
xs	bool	Standout not erased by writing over it (HP 264?)
xt	bool	Tabs are destructive, magic so char (Telaray 1061)

A Sample Entry

The following entry, which describes the Concept-100, is among the more complex entries in the *termcap* file as of this writing. (This particular concept entry is outdated and is used as an example only.)

```
c1|c100|concept100:is=\EU\Ef\E7\E5\E8\EI\ENH\EK\E\200\Eo&\200:\
:al=3*\E^R:am:bs:cd=16*\E^C:ce=16\E^S:cl=2**L:cm=\Ea%+ %+ :co#80:\
:dc=16\E^A:dl=3*\E^B:ei=\E\200:eo:im=\E^P:in:ip=16*:li#24:mi:nd=\E=:
:se=\Ed\Ee:so=\ED\EE:ta=8\t:ul:up=\E:;vb=\Ek\EK:xn:
```

Entries may continue onto multiple lines by giving a \ as the last character of a line, and that empty fields may be included for readability (here between the last field on a line and the first field on the next). Capabilities in *termcap* 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 two letter codes. 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 **co** which indicates the number of columns the terminal has gives the value '80' for the Concept.

Finally, string valued capabilities, such as **cé** (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 after the '=' in such a capability, and padding characters are supplied by the editor after the remainder of the string is sent to provide this delay. The delay can be either an integer, e.g. '20', or an integer 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. When a '**' is specified, it is sometimes useful to give a delay of the form '3.5' specify a delay per unit to tenths of milliseconds.

A number of escape sequences are provided in the string valued capabilities for easy encoding of characters there. A \E maps to an ESCAPE character, ^x maps to a control-x for any appropriate x, and the sequences \n \r \t \b \f give a newline, return, tab, backspace and formfeed. Finally, characters may be given as three octal digits after a \, and the characters ^ and \ may be given as ^ and \\. If it is necessary to place a : in a capability it must be escaped in octal as \072. If it is necessary to place a null character in a

string capability it must be encoded as `\200`. The routines which deal with *termcap* use C strings, and strip the high bits of the output very late so that a `\200` comes out as a `\000` would.

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 *termcap* and to build up a description gradually, using partial descriptions with *ex* to check that they are correct. Be aware that a very unusual terminal may expose deficiencies in the ability of the *termcap* file to describe it or bugs in *ex*. To easily test a new terminal description you can set the environment variable `TERMCAP` to a pathname of a file containing the description you are working on and the editor will look there rather than in *etc/termcap*. `TERMCAP` can also be set to the *termcap* entry itself to avoid reading the file when starting up the editor. (This only works on version 7 systems.)

Basic capabilities

The number of columns on each line for the terminal is given by the `co` numeric capability. If the terminal is a CRT, then the number of lines on the screen is given by the `li` 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, then this is given by the `cl` string capability. If the terminal can backspace, then it should have the `bs` capability, unless a backspace is accomplished by a character other than `^H` (ugh) in which case you should give this character as the `bc` string capability. If it overstrikes (rather than clearing a position when a character is struck over) then it should have the `os` capability.

A very important point here is that the local cursor motions encoded in *termcap* are undefined at the left and top edges of a CRT terminal. The editor will never attempt to backspace around the left edge, nor will it attempt to go up locally off the top. The editor assumes that feeding off the bottom of the screen will cause the screen to scroll up, and the `am` capability tells whether the cursor sticks at the right edge of the screen. If the terminal has switch selectable automatic margins, the *termcap* file usually assumes that this is on, i.e. `am`.

These capabilities suffice to describe hardcopy and glass-tty terminals. Thus the model 33 teletype is described as

```
t3|33|tty33:co#72:os
```

while the Lear Siegler ADM-3 is described as

```
cl|adm3|lsi adm3:am:bs:cl=^Z:li#24:co#80
```

Cursor addressing

Cursor addressing in the terminal is described by a `cm` string capability, with *printf*(3s) like escapes `%x` in it. These substitute to encodings of the current line or column position, while other characters are passed through unchanged. If the `cm` string is thought of as being a function, then its arguments are the line and then the column to which motion is desired, and the `%` encodings have the following meanings:

```
%d    as in printf, 0 origin
%2    like %2d
```

```
%3    like %3d
%.     like %c
%+x    adds x to value, then %.
%>xy  if value > x adds y, no output.
%r     reverses order of line and column, no output
%i     increments line/column (for 1 origin)
%%     gives a single %
%n     exclusive or row and column with 0140 (DM2500)
%B     BCD (16*(x/10)) + (x%10), no output.
%D     Reverse coding (x-2*(x%16)), no output. (Delta Data).
```

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 `cm` capability is `cm=6\E&%r%2c%2Y`. The Microterm ACT-IV needs the current row and column sent preceded by a `^T`, with the row and column simply encoded in binary, `cm=^T%.%..` Terminals which use `%` need to be able to backspace the cursor (`bs` or `bc`), and to move the cursor up one line on the screen (`up` introduced below). This is necessary because it is not always safe to transmit `\t`, `\n` `^D` and `\r`, as the system may change or discard them.

A final example is the LSI ADM-3a, which uses row and column offset by a blank character, thus `cm=\E=%+ %+`.

Cursor motions

If the terminal can move the cursor one position to the right, leaving the character at the current position unchanged, then this sequence should be given as `nd` (non-destructive space). If it can move the cursor up a line on the screen in the same column, this should be given as `up`. If the terminal has no cursor addressing capability, but can home the cursor (to very upper left corner of screen) then this can be given as `ho`; similarly a fast way of getting to the lower left hand corner can be given as `ll`; this may involve going up with `up` from the home position, but the editor will never do this itself (unless `ll` does) because it makes no assumption about the effect of moving up from the home position.

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 `ce`. If the terminal can clear from the current position to the end of the display, then this should be given as `cd`. The editor only uses `cd` from the first column of a line.

Insert/delete line

If the terminal can open a new blank line before the line where the cursor is, this should be given as `al`; 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 `dl`; this is done only from the first position on the line to be deleted. If the terminal can scroll the screen backwards, then this can be given as `sb`, but just `al` suffices. 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 let the editor understand that deleting a line on the screen may bring non-blank lines up from below or that scrolling back with **sb** may bring down non-blank lines.

Insert/delete character

There are two basic kinds of intelligent terminals with respect to insert/delete character which can be described using *termcap*. 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 find out which 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. If your terminal does something different and unusual then you may have to modify the editor to get it to use the insert mode your terminal defines. We have seen no terminals which have an insert mode not falling into one of these two classes.

The editor can handle 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 **lm** the sequence to get into insert mode, or give it an empty value if your terminal uses a sequence to insert a blank position. Give as **ei** the sequence to leave insert mode (give this, with an empty value also if you gave **lm** so). Now give as **ic** any sequence needed to be sent just before sending the character to be inserted. Most terminals with a true insert mode will not give **ic**, terminals which send a sequence to open a screen position should give it here. (Insert mode is preferable to the sequence to open a position on the screen if your terminal has both.) 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**.

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 **mi** to speed up inserting in this case. Omitting **mi** will affect only speed. Some terminals (notably Datamedia's) must not have **mi** because of the way their insert mode works.

Finally, you can specify delete mode by giving **dm** and **ed** to enter and exit delete mode, and **dc** to delete a single character while in delete mode.

Highlighting, underlining, and visible bells

If your terminal has sequences to enter and exit standout mode these can be given as **so** and **se** respectively. If there are several flavors of standout mode (such as inverse video, blinking, or underlining — half bright is not usually an acceptable standout mode unless the terminal is in inverse video mode constantly) the preferred mode is inverse video by itself. 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 **ug** should be given to tell how many spaces are left.

Codes to begin underlining and end underlining can be given as **us** and **ue** 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**. (If the underline code does not move the cursor to the right, give the code followed by a nondestructive space.)

Many 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 newline.

If the terminal has a way of flashing the screen to indicate an error quietly (a bell replacement) then this can be given as **vb**; it must not move the cursor. If the terminal should be placed in a different mode during open and visual modes of *ex*, this can be given as **vs** and **ve**, sent at the start and end of these modes respectively. These can be used to change, e.g., from a underline to a block cursor and back.

If the terminal needs to be in a special mode when running a program that addresses the cursor, the codes to enter and exit this mode can be given as **tl** and **te**. 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.

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 **ks** and **ke**. 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 **k1**, **kr**, **ku**, **kd**, and **kh** respectively. If there are function keys such as **f0**, **f1**, ..., **f9**, the codes they send can be given as **k0**, **k1**, ..., **k9**. If these keys have labels other than the default **f0** through **f9**, the labels can be given as **l0**, **l1**, ..., **l9**. If there are other keys that transmit the same code as the terminal expects for the corresponding

function, such as clear screen, the *termcap* 2 letter codes can be given in the **ko** capability, for example, `:ko=cl,ll,sf,sb:`, which says that the terminal has clear, home down, scroll down, and scroll up keys that transmit the same thing as the `cl`, `ll`, `sf`, and `sb` entries.

The **ma** entry is also used to indicate arrow keys on terminals which have single character arrow keys. It is obsolete but still in use in version 2 of *vi*, which must be run on some minicomputers due to memory limitations. This field is redundant with **kl**, **kr**, **ku**, **kd**, and **kh**. It consists of groups of two characters. In each group, the first character is what an arrow key sends, the second character is the corresponding *vi* command. These commands are **h** for **kl**, **j** for **kd**, **k** for **ku**, **l** for **kr**, and **H** for **kh**. For example, the mime would be `:ma=^Kj^Zk^Xl:` indicating arrow keys left (^H), down (^K), up (^Z), and right (^X). (There is no home key on the mime.)

Miscellaneous

If the terminal requires other than a null (zero) character as a pad, then this can be given as **pc**.

If tabs on the terminal require padding, or if the terminal uses a character other than ^I to tab, then this can be given as **ta**.

Hazeltine terminals, which don't allow `` characters to be printed should indicate **hz**. Datamedia terminals, which echo carriage-return linefeed for carriage return and then ignore a following linefeed should indicate **nc**. Early Concept terminals, which ignore a linefeed immediately after an **am** wrap, should indicate **xn**. If an erase-eol is required to get rid of standout (instead of merely writing on top of it), **xs** should be given. Teleray terminals, where tabs turn all characters moved over to blanks, should indicate **xt**. Other specific terminal problems may be corrected by adding more capabilities of the form **xx**.

Other capabilities include **is**, an initialization string for the terminal, and **if**, the name of a file containing long initialization strings. These strings are expected to properly clear and then set the tabs on the terminal, if the terminal has settable tabs. If both are given, **is** will be printed before **if**. This is useful where **if** is `/usr/lib/tabset/std` but **is** clears the tabs first.

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 **tc** can be given with the name of the similar terminal. This capability must be *last* and the combined length of the two entries must not exceed 1024. Since *termlib* routines search the entry from left to right, and since the **tc** capability is replaced by the corresponding entry, the capabilities given at the left override the ones in the similar terminal. A capability can be cancelled with **xx@** where **xx** is the capability. For example, the entry

```
hn|2621nl:ks@:ke@:tc=2621:
```

defines a `2621nl` that does not have the **ks** or **ke** 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

`/etc/termcap` file containing terminal descriptions

SEE ALSO

`ex(1)`, `more(1)`, `tset(1)`, `ul(1)`, `vi(1)`, `termcap(3)`.

BUGS

Ex allows only 256 characters for string capabilities, and the routines in *termcap(3)* do not check for overflow of this buffer. The total length of a single entry (excluding only escaped newlines) may not exceed 1024.

The **ma**, **vs**, and **ve** entries are specific to the *vi* program.

Not all programs support all entries. There are entries that are not supported by any program.

AUTHOR

William Joy

Mark Horton added underlining and keypad support

NAME

types — primitive system data types

SYNOPSIS

#include <sys/types.h>

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)*.

NAME

udp - Internet User Datagram Protocol

SYNOPSIS

```
#include <net/socket.h>
#include <net/in.h>
```

```
struct sockproto proto = { PF_INET, IPPROTO_UDP };
socket(SOCK_DGRAM, &proto, address, options);
struct sockaddr_in *address; int options;
```

DESCRIPTION

UDP is a simple, unreliable datagram protocol which is used to support the `SOCK_DGRAM` abstraction for the Internet protocol family. *UDP* sockets are connectionless, and are normally used with the *send(2)* and *receive(2)* calls, though the *connect(2)* call may also be used to fix the destination for future packets (in which case the *read(2)* and *write(2)* system calls may be used).

UDP address formats are identical to those used by *TCP*. In particular *UDP* provides a port identifier in addition to the normal Internet address format. Note that the *UDP* port space is separate from the *TCP* port space (i.e., a *UDP* port may not be "connected" to a *TCP* port). In addition broadcast packets may be sent (assuming the underlying network supports this) by using a reserved "broadcast address"; this address is network interface dependent.

DIAGNOSTICS

`EISCONN` when trying to establish a connection on a socket which already has one, or when trying to send a datagram with the destination address specified and the socket is already connected;

`ENOTCONN` when trying to send a datagram, but no destination address is specified, and the socket hasn't been connected;

`ENOBUFS` when the system runs out of memory for an internal data structure;

`EADDRINUSE` when an attempt is made to create a socket with a port which has already been allocated;

`EADDRNOTAVAIL` when an attempt is made to create a socket with a network address for which no network interface exists.

SEE ALSO

`inet(5)`, `net(5)`.

NAME

intro — introduction to games

DESCRIPTION

This section describes the recreational and educational programs found in the directory `/usr/games`. The availability of these programs may vary from system to system.

NAME

adventure — an exploration game

SYNOPSIS

`/usr/games/adventure`

DESCRIPTION

The object of the game is to locate and explore Colossal Cave, find the treasures hidden there, and bring them back to the building with you. The program is self-describing to a point, but part of the game is to discover its rules.

To terminate a game, type "quit"; to save a game for later resumption, type "suspend".

BUGS

Saving a game creates a large executable file instead of just the information needed to resume the game.

NAME

aliens — The alien invaders attack the earth

SYNOPSIS

/usr/games/aliens

DESCRIPTION

This is a UNIX version of Space Invaders. The program is pretty much self documenting.

FILES

/usr/games/lib/aliens.log Score file

BUGS

The program is a CPU hog. It needs to be re-written. It doesn't do well on terminals that run slower than 9600 baud.

NAME

arithmetic — provide drill in number facts

SYNOPSIS

/usr/games/arithmetic [+ - x/] [range]

DESCRIPTION

Arithmetic types out simple arithmetic problems, and waits for an answer to be typed in. If the answer is correct, it types back "Right!", and a new problem. If the answer is wrong, it replies "What?", and waits for another answer. Every twenty problems, it publishes statistics on correctness and the time required to answer.

To quit the program, type an interrupt (delete).

The first optional argument determines the kind of problem to be generated; +, -, x, and / respectively cause addition, subtraction, multiplication, and division problems to be generated. One or more characters can be given; if more than one is given, the different types of problems will be mixed in random order; default is + -.

Range is a decimal number; all addends, subtrahends, differences, multipliers, divisors, and quotients will be less than or equal to the value of *range*. Default *range* is 10.

At the start, all numbers less than or equal to *range* are equally likely to appear. If the respondent makes a mistake, the numbers in the problem which was missed become more likely to reappear.

As a matter of educational philosophy, the program will not give correct answers, since the learner should, in principle, be able to calculate them. Thus the program is intended to provide drill for someone just past the first learning stage, not to teach number facts *de novo*. For almost all users, the relevant statistic should be time per problem, not percent correct.

NAME

autorobots — Escape from the automatic robots

SYNOPSIS

`/usr/games/autorobots`

DESCRIPTION

The object of the game *autorobots* is to move around inside of the box on the screen without getting eaten by the robots chasing you and without running into any robots or junk heaps. The robots move continuously.

If a robot runs into another robot or junk heap while chasing you, they crash and leave a junk heap.

You start out with 10 robots worth 10 points each. If you defeat all of them, you get 20 robots worth 20 points each. Then 30, etc. Until you get eaten!

The game keeps track of the top ten scores and prints them at the end of the game.

The valid commands are described on the screen.

NAME

back — the game of backgammon

SYNOPSIS

`/usr/games/back`

DESCRIPTION

Back is a program which provides a partner for the game of backgammon. It is designed to play at three different levels of skill, one of which you must select. In addition to selecting the opponent's level, you may also indicate that you would like to roll your own dice during your turns (for the superstitious players). You will also be given the opportunity to move first. The practice of each player rolling one die for the first move is not incorporated.

The points are numbered 1–24, with 1 being white's extreme inner table, 24 being brown's inner table, 0 being the bar for removed white pieces and 25 the bar for brown. For details on how moves are expressed, type *y* when *back* asks "Instructions?" at the beginning of the game. When *back* first asks "Move?", type ? to see a list of move options other than entering your numerical move.

When the game is finished, *back* will ask you if you want the log. If you respond with *y*, *back* will attempt to append to or create a file **back.log** in the current directory.

FILES

<code>/usr/games/lib/backrules</code>	rules file
<code>/tmp/b*</code>	log temp file
<code>back.log</code>	log file

BUGS

The only level really worth playing is "expert", and it only plays the forward game.

Back will complain loudly if you attempt to make too *many* moves in a turn, but will become very silent if you make too *few*.

Doubling is not implemented.

NAME

bcd — convert to antique media

SYNOPSIS

`/usr/games/bcd` text

DESCRIPTION

Bcd converts the literal *text* into a form familiar to old-timers.

This program works best on hard copy terminals.

NAME

bj — the game of black jack

SYNOPSIS

/usr/games/bj

DESCRIPTION

Bj is a serious attempt at simulating the dealer in the game of black jack (or twenty-one) as might be found in Reno. The following rules apply:

The bet is \$2 every hand.

A player "natural" (black jack) pays \$3. A dealer natural loses \$2. Both dealer and player naturals is a "push" (no money exchange).

If the dealer has an ace up, the player is allowed to make an "insurance" bet against the chance of a dealer natural. If this bet is not taken, play resumes as normal. If the bet is taken, it is a side bet where the player wins \$2 if the dealer has a natural and loses \$1 if the dealer does not.

If the player is dealt two cards of the same value, he is allowed to "double". He is allowed to play two hands, each with one of these cards. (The bet is doubled also; \$2 on each hand.)

If a dealt hand has a total of ten or eleven, the player may "double down". He may double the bet (\$2 to \$4) and receive exactly one more card on that hand.

Under normal play, the player may "hit" (draw a card) as long as his total is not over twenty-one. If the player "busts" (goes over twenty-one), the dealer wins the bet.

When the player "stands" (decides not to hit), the dealer hits until he attains a total of seventeen or more. If the dealer busts, the player wins the bet.

If both player and dealer stand, the one with the largest total wins. A tie is a push.

The machine deals and keeps score. The following questions will be asked at appropriate times. Each question is answered by y followed by a new-line for "yes", or just new-line for "no".

? (means, "do you want a hit?")

Insurance?

Double down?

Every time the deck is shuffled, the dealer so states and the "action" (total bet) and "standing" (total won or lost) is printed. To exit, hit the interrupt key (DEL) and the action and standing will be printed.

NAME

chase — Try to escape the killer robots

SYNOPSIS

`/usr/games/chase [nrobots] [nfences]`

DESCRIPTION

The object of the game *chase* is to move around inside of the box on the screen without getting eaten by the robots chasing you and without running into anything.

If a robot runs into another robot while chasing you, they crash and leave a junk heap. If a robot runs into a fence, it is destroyed.

If you can survive until all the robots are destroyed, you have won!

If you do not specify either *nrobots* or *nfences*, chase will prompt you for them.

The valid commands are described on the screen.

NAME

craps — the game of craps

SYNOPSIS

`/usr/games/craps`

DESCRIPTION

Craps is a form of the game of craps that is played in Las Vegas. The program simulates the *roller*, while the user (the *player*) places bets. The player may choose, at any time, to bet with the roller or with the *House*. A bet of a negative amount is taken as a bet with the House, any other bet is a bet with the roller.

The player starts off with a "bankroll" of \$2,000.

The program prompts with:

bet?

The bet can be all or part of the player's bankroll. Any bet over the total bankroll is rejected and the program prompts with **bet?** until a proper bet is made.

Once the bet is accepted, the roller throws the dice. The following rules apply (the player wins or loses depending on whether the bet is placed with the roller or with the House; the odds are even). The *first* roll is the roll immediately following a bet:

1. On the first roll:

7 or 11	wins for the roller;
2, 3, or 12	wins for the House;
any other number	is the <i>point</i> , roll again (Rule 2 applies).

2. On subsequent rolls:

point	roller wins;
7	House wins;
any other number	roll again.

If a player loses the entire bankroll, the House will offer to lend the player an additional \$2,000. The program will prompt:

marker?

A **yes** (or **y**) consummates the loan. Any other reply terminates the game.

If a player owes the House money, the House reminds the player, before a bet is placed, how many markers are outstanding.

If, at any time, the bankroll of a player who has outstanding markers exceeds \$2,000, the House asks:

Repay marker?

A reply of **yes** (or **y**) indicates the player's willingness to repay the loan. If only 1 marker is outstanding, it is immediately repaid. However, if more than 1 marker are outstanding, the House asks:

How many?

markers the player would like to repay. If an invalid number is entered (or just a carriage return), an appropriate message is printed and the program will prompt with **How many?** until a valid number is entered.

If a player accumulates 10 markers (a total of \$20,000 borrowed from the House), the program informs the player of the situation and exits.

Should the bankroll of a player who has outstanding markers exceed \$50,000, the *total* amount of money borrowed will be *automatically* repaid to the House.

Any player who accumulates \$100,000 or more breaks the bank. The program then prompts:

New game?

to give the House a chance to win back its money.

Any reply other than **yes** is considered to be a **no** (except in the case of **bet?** or **How many?**). To exit, send an interrupt (break), DEL, or control-D. The program will indicate whether the player won, lost, or broke even.

MISCELLANEOUS

The random number generator for the die numbers uses the seconds from the time of day. Depending on system usage, these numbers, at times, may seem strange but occurrences of this type in a real dice situation are not uncommon.

NAME

cribbage — the card game cribbage

SYNOPSIS

`/usr/games/cribbage [-[r][e][q]] name ...`

DESCRIPTION

Cribbage plays the card game cribbage, with the program playing one hand and the user the other. The program will initially ask the user if the rules of the game are needed -- if so, it will print out the appropriate section from *According to Hoyle* with *more (1)*.

Cribbage options include:

- e When the player makes a mistake scoring his hand or crib, provide an explanation of the correct score. (This is especially useful for beginning players.)
- q Print a shorter form of all messages -- this is only recommended for users who have played the game without specifying this option.
- r Instead of asking the player to cut the deck, the program will randomly cut the deck.

Cribbage first asks the player whether he wishes to play a short game (once around, to 61) or a long game (twice around, to 121). A response of 's' will result in a short game, any other response will play a long game.

At the start of the first game, the program asks the player to cut the deck to determine who gets the first crib. The user should respond with a number between 0 and 51, indicating how many cards down the deck is to be cut. The player who cuts the lower ranked card gets the first crib. If more than one game is played, the loser of the previous game gets the first crib in the current game.

For each hand, the program first prints the player's hand, whose crib it is, and then asks the player to discard two cards into the crib. The cards are prompted for one per line, and are typed as explained below.

After discarding, the program cuts the deck (if it is the player's crib) or asks the player to cut the deck (if it's its crib); in the later case, the appropriate response is a number from 0 to 39 indicating how far down the remaining 40 cards are to be cut.

After cutting the deck, play starts with the non-dealer (the person who doesn't have the crib) leading the first card. Play continues, as per cribbage, until all cards are exhausted. The program keeps track of the scoring of all points and the total of the cards on the table.

After play, the hands are scored. The program requests the player to score his hand (and the crib, if it is his) by printing out the appropriate cards (and the cut card enclosed in brackets). Play continues until one player reaches the game limit (61 or 121).

A carriage return when a numeric input is expected is equivalent to typing the lowest legal value; when cutting the deck this is equivalent to choosing the top card.

Cards are specified as rank followed by suit. The ranks may be specified as one of: 'a', '2', '3', '4', '5', '6', '7', '8', '9', 't', 'j', 'q', and 'k', or alternatively, one of: ace, two, three, four, five, six, seven, eight, nine, ten, jack, queen, and king. Suits may be specified as: 's', 'h', 'd', and 'c', or alternatively as: spades, hearts, diamonds, and clubs. A card may be specified as: <rank> <suit>, or: <rank> of <suit>. If the single letter rank and suit designations are used, the space separating the suit and rank may be left out. Also, if only one card of the desired rank is playable, typing the rank is sufficient. For example, if your hand was 2H, 4D, 5C, 6H, JC, KD and it was desired to discard the king of diamonds, any of the following could be typed: k, king, kd, k d, k of d, king d, king of d, k diamonds, k of diamonds, king diamonds, or king of diamonds.

FILES

/usr/games/cribbage

AUTHOR

Earl T. Cohen

NAME

fish - play "Go Fish"

SYNOPSIS

/usr/games/fish

DESCRIPTION

Fish plays the game of Go Fish, a childrens' card game. The Object is to accumulate 'books' of 4 cards with the same face value. The players alternate turns; each turn begins with one player selecting a card from his hand, and asking the other player for all cards of that face value. If the other player has one or more cards of that face value in his hand, he gives them to the first player, and the first player makes another request. Eventually, the first player asks for a card which is not in the second player's hand: he replies 'GO FISH!' The first player then draws a card from the 'pool' of undealt cards. If this is the card he had last requested, he draws again. When a book is made, either through drawing or requesting, the cards are laid down and no further action takes place with that face value.

To play the computer, simply make guesses by typing a, 2, 3, 4, 5, 6, 7, 8, 9, 10, j, q, or k when asked. Hitting return gives you information about the size of my hand and the pool, and tells you about my books. Saying 'p' as a first guess puts you into 'pro' level; the default is pretty dumb.

NAME

fortune — print a random, hopefully interesting, adage

SYNOPSIS

fortune

DESCRIPTION

Fortune prints out a random adage.

FILES

/usr/games/lib/fortunes

NAME

hangman -- guess the word

SYNOPSIS

/usr/games/hangman [arg]

DESCRIPTION

Hangman chooses a word at least seven letters long from a dictionary. The user is to guess letters one at a time.

The optional argument *arg* names an alternate dictionary.

FILES

/usr/lib/w2006

BUGS

Hyphenated compounds are run together.

NAME

life — play the game of life

SYNOPSIS

life [-r]

DESCRIPTION

Life is a pattern generating game set up for interactive use on a video terminal. The way it operates is: You use a series of commands to set up a pattern on the screen then let it generate further patterns from that pattern.

The algorithm used is: For each square in the matrix, look at it and its eight adjacent neighbors. If the present square is not occupied and exactly three of its neighbor squares are occupied, then that square will be occupied in the next pattern. If the present square is occupied and two or three of its neighbor squares are occupied, then that square will be occupied in the next pattern. Otherwise, the present square will not be occupied in the next pattern.

The edges of the screen are normally treated as an unoccupied void. If you specify the *-r* option on the command line, the screen is treated as a sphere; that is, the top and bottom lines are considered adjacent and the left and right columns are considered adjacent.

The pattern generation number and the number of occupied squares are displayed in the lower left hand corner.

Below is a list of commands available to the user. A # stands for any number. A ^ followed by a capital letter represents a control character.

- #, #a** Add a block of elements. The first number specifies the horizontal width. The second number specifies the vertical width. If a number is not specified, the default is 1.
- #c** Step through the next # patterns. If no number is specified, step forever. The operation can be aborted by typing rubout (delete).
- #, #d** Delete a block of elements. The first number specifies the horizontal width. The second number specifies the vertical width. If a number is not specified, the default is 1.
- #f** Generate a little flier at the present location. The number (modulo 8) determines the direction.
- #, #g** Move to absolute screen location. The first number specifies the horizontal location. The second number specifies the vertical location. If a number is not specified, the default is 0.
- #h** Move left # steps. If no number is specified, the default is 1.
- #j** Move down # steps. The default is 1.
- #k** Move up # steps. The default is 1.
- #l** Move right # steps. The default is 1.
- #n** Step through the next # patterns. If no number is specified, generate the next pattern. The operation can be aborted by typing rubout (delete).
- p** Put the last yanked or deleted block at the present location.

LIFE (6)

(UniSoft)

LIFE (6)

q Quit.

#, #y Yank a block of elements. The first number specifies the horizontal width. The second number specifies the vertical width. If a number is not specified, the default is 1.

C Clear the pattern.

#F Generate a big flier at the present location. The number (modulo 8) determines the direction.

#H Move to the left margin.

#J Move to the bottom margin.

#K Move to the top margin.

#L Move to the right margin.

#^H Move left # steps. If no number is specified, the default is 1.

#^J Move down # steps. The default is 1.

#^K Move up # steps. The default is 1.

#^L Move right # steps. The default is 1.

^R Redraw the screen. This is used for those occasions when the terminal screws up.

. Repeat the last add (a) or delete (d) operation.

; Repeat the last move (h, j, k, l) operation.

BUGS

The following features are planned but not implemented:

#, #S Save the selected area in a file.

R Restore from a file.

m Generate a macro command.

! Shell escape.

e Edit a file.

i Input commands from a file.

AUTHOR

Asa Romberger

MAZE (6)**MAZE (6)****NAME**

maze - generate a maze

SYNOPSIS

/usr/games/maze

DESCRIPTION

Maze asks a few questions and then prints a maze.

BUGS

Some mazes (especially small ones) have no solutions.

NAME

moo -- guessing game

SYNOPSIS`/usr/games/moo`**DESCRIPTION**

Moo is a guessing game imported from England. The computer picks a number consisting of four distinct decimal digits. The player guesses four distinct digits being scored on each guess. A "cow" is a correct digit in an incorrect position. A "bull" is a correct digit in a correct position. The game continues until the player guesses the number (a score of four bulls).

NAME

number – convert Arabic numerals to English

SYNOPSIS

/usr/games/number

DESCRIPTION

Number copies the standard input to the standard output, changing each decimal number to a fully spelled out version.

NAME

quiz — test your knowledge

SYNOPSIS

/usr/games/quiz [-i file] [-t] [category1 category2]

DESCRIPTION

Quiz gives associative knowledge tests on various subjects. It asks items chosen from *category1* and expects answers from *category2*, or vice versa. If no categories are specified, *quiz* gives instructions and lists the available categories.

Quiz tells a correct answer whenever you type a bare new-line. At the end of input, upon interrupt, or when questions run out, *quiz* reports a score and terminates.

The **-t** flag specifies “tutorial” mode, where missed questions are repeated later, and material is gradually introduced as you learn.

The **-i** flag causes the named file to be substituted for the default index file. The lines of these files have the syntax:

```

line      = category new-line | category : line
category  = alternate | category | alternate
alternate = empty | alternate primary
primary   = character | | category | | option
option    = { category }
```

The first category on each line of an index file names an information file. The remaining categories specify the order and contents of the data in each line of the information file. Information files have the same syntax. Backslash `\` is used as with *sh*(1) to quote syntactically significant characters or to insert transparent new-lines into a line. When either a question or its answer is empty, *quiz* will refrain from asking it.

FILES

```

/usr/games/lib/quiz/index
/usr/games/lib/quiz/*
```

BUGS

The construct “a|ab” doesn’t work in an information file. Use “a{b}”.

NAME

rain — animated raindrops display

SYNOPSIS

rain

DESCRIPTION

Rain's display is modeled after the VAX/VMS program of the same name. The terminal has to be set for 9600 baud to obtain the proper effect.

As with all programs that use *termcap*, the TERM environment variable must be set (and exported) to the type of the terminal being used.

FILES

/etc/termcap

AUTHOR

Eric P. Scott

NAME

robots — Escape from the robots

SYNOPSIS

/usr/games/robots

DESCRIPTION

The object of the game *robots* is to move around inside of the box on the screen without getting eaten by the robots chasing you and without running into anything.

If a robot runs into another robot while chasing you, they crash and leave a junk heap.

You start out with 10 robots worth 10 points each. If you defeat all of them, you get 20 robots worth 20 points each. Then 30, etc. Until you get eaten!

The game keeps track of the top ten scores and prints them at the end of the game.

The valid commands are described on the screen.

NAME

trek - trekkie game

SYNOPSIS

/usr/games/trek [[-a] file]

DESCRIPTION

Trek is a game of space glory and war. Below is a summary of commands. For complete documentation, see *Trek* by Eric Allman.

If a filename is given, a log of the game is written onto that file. If the `-a` flag is given before the filename, that file is appended to, not truncated.

The game will ask you what length game you would like. Valid responses are short, medium, and long. You may also type restart, which restarts a previously saved game. You will then be prompted for the skill, to which you must respond novice, fair, good, expert, commadore, or impossible. You should normally start out with a novice and work up.

In general, throughout the game, if you forget what is appropriate the game will tell you what it expects if you just type in a question mark.

COMMAND SUMMARY

abandon	capture
cloak up/down	
computer request; ...	damages
destruct	dock
help	impulse course distance
lrscan	move course distance
phasers automatic amount	
phasers manual amt1 course1 spread1 ...	
torpedo course [yes] angle/no	
ram course distance	rest time
shell	shields up/down
srscan [yes/no]	
status	terminate yes/no
undock	visual course
warp warp_factor	

AUTHOR

Eric Allman

NAME

ttt, cubic - tic-tac-toe

SYNOPSIS

/usr/games/ttt
/usr/games/cubic

DESCRIPTION

Ttt is the X and O game popular in the first grade. This is a learning program that never makes the same mistake twice.

Although it learns, it learns slowly. It must lose nearly 80 games to completely know the game.

Cubic plays three-dimensional tic-tac-toe on a 4×4×4 board. Moves are specified as a sequence of three coordinate numbers in the range 1-4.

FILES

/usr/games/ttt.k learning file

NAME

twinkle -- twinkle stars on the screen

SYNOPSIS

/usr/games/twinkle [-+ [s save]] [density1] [density2]

DESCRIPTION

Twinkle causes a specified density of 'stars' to twinkle on the screen. The following options are available;

- print out the present screen density (the percentage of the screen that will be filled with stars) in the lower left hand corner of the screen. This number will change as stars go on and off.
- + do not 'randomize' before starting. The screen starts out completely blank and stars are added, bit by bit. In this case the density rises beyond the specified density, then falls to the required percentage.
- s save binary density on file 'save', in case you want to see the density curve that a particular density specification produced during the life of the show.

density If no density is specified, density is .5 (50% of the screen will be filled with stars).

If only *density1* is given, density is $1/\text{density1}$

If both *density1* and *density2* are given, *density* is the resultant of $\text{density1}/\text{density1} + \text{density2}$.

EXAMPLE

twinkle -+ 2 6

would start from a blank screen and twinkle stars to a final density of 2/8, or 25%. The densities would be shown in the lower left hand corner, as a three-place decimal.

AUTHOR

Asa Romberger

NAME

worm — Play the growing worm game

SYNOPSIS

worm [*size*]

DESCRIPTION

In *worm*, you are a little worm, your body is the "o"'s on the screen and your head is the "@". You move with the hjkl keys (as in the game snake). If you don't press any keys, you continue in the direction you last moved. The upper case HJKL keys move you as if you had pressed several (9 for HL and 5 for JK) of the corresponding lower case key (unless you run into a digit, then it stops).

On the screen you will see a digit; if your worm eats the digit, it will grow longer. The actual amount by which the worm will grow longer depends upon which digit was eaten. The object of the game is to see how long you can make the worm grow.

The game ends when the worm runs into either the sides of the screen, or itself. The current score (how much the worm has grown) is kept in the upper left corner of the screen.

The optional argument, if present, is the initial length of the worm.

BUGS

If the initial length of the worm is set to less than one or more than 75, various strange things happen.

NAME

worms - animate worms on a display terminal

SYNOPSIS

worms [-field] [-length #] [-number #] [-trail]

DESCRIPTION

-field makes a "field" for the worm(s) to eat; -trail causes each worm to leave a trail behind it. You can figure out the rest by yourself.

FILES

/etc/termcap

DIAGNOSTICS

Invalid length

Value not in range $2 \leq \text{length} \leq 1024$

Invalid number of worms

Value not in range $1 \leq \text{number} \leq 40$

TERM: parameter not set

The **TERM** environment variable is not defined. Do

TERM=terminal type

export TERM

Unknown terminal type

Your terminal type (as determined from the **TERM** environment variable) is not defined in **/etc/termcap**.

Terminal not capable of cursor motion

Your terminal is too stupid to run this program.

Out of memory

This should never happen.

BUGS

The lower-right-hand character position will not be updated properly on a terminal that wraps at the right margin.

Terminal initialization is not performed.

AUTHOR

Eric P. Scott

NAME

wump — the game of hunt-the-wumpus

SYNOPSIS

/usr/games/wump

DESCRIPTION

Wump plays the game of " Hunt the Wumpus." A Wumpus is a creature that lives in a cave with several rooms connected by tunnels. You wander among the rooms, trying to shoot the Wumpus with an arrow, meanwhile avoiding being eaten by the Wumpus and falling into Bottomless Pits. There are also Super Bats which are likely to pick you up and drop you in some random room.

The program asks various questions which you answer one per line; it will give a more detailed description if you want.

This program is based on one described in *People's Computer Company*, 2, 2 (November 1973).

BUGS

It will never replace Adventure.