

NAME

ar – archive and library maintainer

SYNOPSIS

ar key afile name ...

DESCRIPTION

Ar maintains groups of files combined into a single archive file. Its main use is to create and update library files as used by the loader. It can be used, though, for any similar purpose.

Key is one character from the set **drtux**, optionally concatenated with **v**. *Afile* is the archive file. The *names* are constituent files in the archive file. The meanings of the *key* characters are:

d means delete the named files from the archive file.

r means replace the named files in the archive file. If the archive file does not exist, **r** will create it. If the named files are not in the archive file, they are appended.

t prints a table of contents of the archive file. If no names are given, all files in the archive are tabled. If names are given, only those files are tabled.

u is similar to **r** except that only those files that have been modified are replaced. If no names are given, all files in the archive that have been modified will be replaced by the modified version.

x will extract the named files. If no names are given, all files in the archive are extracted. In neither case does **x** alter the archive file.

v means verbose. Under the verbose option, *ar* gives a file-by-file description of the making of a new archive file from the old archive and the constituent files. The following abbreviations are used:

- c** copy
- a** append
- d** delete
- r** replace
- x** extract

FILES

/tmp/vtm? temporary

SEE ALSO

ld (I), archive (V)

BUGS

Option **tv** should be implemented as a table with more information.

There should be a way to specify the placement of a new file in an archive. Currently, it is placed at the end.

Since *ar* has not been rewritten to deal properly with the new file system modes, extracted files have mode 666.

For the same reason, names are only maintained to 8 characters.

NAME

as – assembler

SYNOPSIS

as [-] name ...

DESCRIPTION

As assembles the concatenation of the named files. If the optional first argument - is used, all undefined symbols in the assembly are treated as global.

The output of the assembly is left on the file **a.out**. It is executable if no errors occurred during the assembly, and if there were no unresolved external references.

FILES

/etc/as2	pass 2 of the assembler
/tmp/atm[1-4]?	temporary
a.out	object

SEE ALSO

ld(I), nm(I), db(I), a.out(V), 'UNIX Assembler Manual'.

DIAGNOSTICS

When an input file cannot be read, its name followed by a question mark is typed and assembly ceases. When syntactic or semantic errors occur, a single-character diagnostic is typed out together with the line number and the file name in which it occurred. Errors in pass 1 cause cancellation of pass 2. The possible errors are:

)	Parentheses error
]	Parentheses error
<	String not terminated properly
*	Indirection used illegally
.	Illegal assignment to '.'
A	Error in address
B	Branch instruction is odd or too remote
E	Error in expression
F	Error in local ('f' or 'b') type symbol
G	Garbage (unknown) character
I	End of file inside an if
M	Multiply defined symbol as label
O	Word quantity assembled at odd address
P	'.' different in pass 1 and 2
R	Relocation error
U	Undefined symbol
X	Syntax error

BUGS

Symbol table overflow is not checked. **x** errors can cause incorrect line numbers in following diagnostics.

NAME

cat – concatenate and print

SYNOPSIS

cat file ...

DESCRIPTION

Cat reads each file in sequence and writes it on the standard output. Thus:

cat file

is about the easiest way to print a file. Also:

cat file1 file2 >file3

is about the easiest way to concatenate files.

If no input file is given *cat* reads from the standard input file.

If the argument – is encountered, *cat* reads from the standard input file.

SEE ALSO

pr(I), cp(I)

DIAGNOSTICS

none; if a file cannot be found it is ignored.

BUGS

cat x y >x and **cat x y >y** cause strange results.

NAME

`cc` – C compiler

SYNOPSIS

`cc` [`-c`] [`-p`] [`-O`] [`-S`] [`-P`] file ...

DESCRIPTION

`Cc` is the UNIX C compiler. It accepts three types of arguments:

Arguments whose names end with `‘.c’` are taken to be C source programs; they are compiled, and each object program is left on the file whose name is that of the source with `‘.o’` substituted for `‘.c’`. The `‘.o’` file is normally deleted, however, if a single C program is compiled and loaded all at one go.

The following flags are interpreted by `cc`. See *ld (I)* for load-time flags.

- `-c` Suppress the loading phase of the compilation, and force an object file to be produced even if only one program is compiled.
- `-p` Arrange for the compiler to produce code which counts the number of times each routine is called; also, if loading takes place, replace the standard startup routine by one which automatically calls the *monitor* subroutine (III) at the start and arranges to write out a *mon.out* file at normal termination of execution of the object program. An execution profile can then be generated by use of *prof (I)*.
- `-O` Invoke the experimental object-code optimizer.
- `-S` Compile the named C programs, and leave the assembler-language output on corresponding files suffixed `‘.s’`.
- `-P` Run only the macro preprocessor on the named C programs, and leave the output on corresponding files suffixed `‘.i’`.

Other arguments are taken to be either loader flag arguments, or C-compatible object programs, typically produced by an earlier `cc` run, or perhaps libraries of C-compatible routines. These programs, together with the results of any compilations specified, are loaded (in the order given) to produce an executable program with name **a.out**.

FILES

file.c	input file
file.o	object file
a.out	loaded output
/tmp/ctm?	temporary
/lib/c[01]	compiler
/lib/c2	optional optimizer
/lib/crt0.o	runtime startoff
/lib/mcrt0.o	runtime startoff of monitoring
/lib/libc.a	builtin functions, etc.
/lib/liba.a	system library

SEE ALSO

“Programming in C— a tutorial,” C Reference Manual, *monitor (III)*, *prof (I)*, *cdb (I)*, *ld (I)*.

DIAGNOSTICS

The diagnostics produced by C itself are intended to be self-explanatory. Occasional messages may be produced by the assembler or loader. Of these, the most mystifying are from the assembler, in particular `‘m,’` which means a multiply-defined external symbol (function or data).

BUGS

NAME

cdb – C debugger

SYNOPSIS

cdb [core [a.out]]

DESCRIPTION

Cdb is a debugger for use with C programs. The first argument is a core-image file; if not given, “core” is used. The second argument is the object program (containing a symbol table); if not given “a.out” is used. An acceptable core and object file must both be present.

Commands to *cdb* consist of an address, followed by a single command character. If no address is given the last-printed address is used. An address may be followed by a comma and a number, in which case the command applies to the appropriate number of successive addresses.

Addresses are expressions composed of names, decimal numbers, and octal numbers (which begin with ‘0’) and separated by ‘+’ and ‘-’. Evaluation proceeds left-to-right. The construction “name[expression]” assumes that *name* is a pointer to an integer and is equivalent to the contents of the named cell plus twice the expression.

The command characters are:

- / print the addressed words in octal.
- = print the value of the addressed expression.
- ’ print the addressed bytes as characters.
- " take the address as a pointer to a sequence of characters, and print the characters up to a null byte.
- & If there is any symbol which has the same value as the address, print the symbol’s name.
- \$ print a stack trace of the terminated program. The calls are listed in the order made; the actual arguments to each routine are given in octal.

SEE ALSO

cc (I), db (I), C Reference Manual

BUGS

It’s still very feeble, even compared with *db* (I). The stack trace is also pretty vulnerable to corruption, and often doesn’t work.

NAME

chdir – change working directory

SYNOPSIS

chdir directory

DESCRIPTION

Directory becomes the new working directory. The process must have execute (search) permission in *directory*.

Because a new process is created to execute each command, *chdir* would be ineffective if it were written as a normal command. It is therefore recognized and executed by the Shell.

SEE ALSO

sh(I)

BUGS

NAME

chmod – change mode

SYNOPSIS

chmod octal file ...

DESCRIPTION

The octal mode replaces the mode of each of the files. The mode is constructed from the OR of the following modes:

- 4000 set user ID on execution
- 2000 set group ID on execution
- 0400 read by owner
- 0200 write by owner
- 0100 execute (search in directory) by owner
- 0070 read, write, execute (search) by group
- 0007 read, write, execute (search) by others

Only the owner of a file (or the super-user) may change its mode.

SEE ALSO

ls(I)

BUGS

NAME

chown – change owner

SYNOPSIS

chown owner file ...

DESCRIPTION

Owner becomes the new owner of the files. The owner may be either a decimal UID or a login name found in the password file.

Only the owner of a file (or the super-user) is allowed to change the owner. Unless it is done by the super-user, the set-user-ID permission bit is turned off as the owner of a file is changed.

FILES

/etc/passwd

BUGS

NAME

`cmp` – compare two files

SYNOPSIS

`cmp` file1 file2

DESCRIPTION

The two files are compared for identical contents. Discrepancies are noted by giving the offset and the differing words, all in octal.

SEE ALSO

`proof` (I), `comm` (I)

BUGS

If the shorter of the two files is of odd length, *cmp* acts as if a null byte had been appended to it. The *offset* is only a single-precision number.

NAME

comm – print lines common to two files

SYNOPSIS

comm [- [**123**]] file1 file2

DESCRIPTION

Comm reads *file1* and *file2*, which should be in sort, and produces a three column output: lines only in *file1*; lines only in *file2*; and lines in both files.

The output is written on the standard output.

Flags 1, 2, or 3 suppress printing of the corresponding column. Thus **comm -12** prints only the lines common to the two files; **comm -23** prints only lines in the first file but not in the second; **comm -123** is a no-op.

SEE ALSO

cmp (I), diff (I)

BUGS

NAME

cp – copy

SYNOPSIS

cp [**-t**] file1 file2

DESCRIPTION

The first file is copied onto the second. The mode and owner of the target file are preserved if it already existed; the mode of the source file is used otherwise.

If *file2* is a directory, then the target file is a file in that directory with the file-name of *file1*.

No one is quite sure what the flag **-t** does.

SEE ALSO

cat (I), pr (I), mv (I)

BUGS

Copying a file onto itself destroys its contents.

NAME

cref – make cross reference listing

SYNOPSIS

cref [**-acilostux123**] name ...

DESCRIPTION

Cref makes a cross reference listing of program files in assembler or C format. The files named as arguments in the command line are searched for symbols in the appropriate syntax.

The output report is in four columns:

(1)	(2)	(3)	(4)
symbol	file	see	text as it appears in file
		below	

Cref uses either an *ignore* file or an *only* file. If the **-i** option is given, it will take the next available argument to be an *ignore* file name; if the **-o** option is given, the next available argument will be taken as an *only* file name. *Ignore* and *only* files should be lists of symbols separated by new lines. If an *ignore* file is given, all the symbols in that file will be ignored in columns (1) and (3) of the output. If an *only* file is given, only symbols appearing in that file will appear in column (1). Only one of the options **-i** or **-o** may be used. The default setting is **-i**. Assembler predefined symbols or C keywords are ignored.

The **-s** option causes current symbols to be put in column 3. In the assembler, the current symbol is the most recent name symbol; in C, the current function name. The **-l** option causes the line number within the file to be put in column 3.

The **-t** option causes the next available argument to be used as the name of the intermediate temporary file (instead of /tmp/crt??). The file is created and is not removed at the end of the process.

Options:

- a** assembler format (default)
- c** C format input
- i** use *ignore* file (see above)
- l** put line number in col. 3 (instead of current symbol)
- o** use *only* file (see above)
- s** current symbol in col. 3 (default)
- t** user supplied temporary file
- u** print only symbols that occur exactly once
- x** print only C external symbols
- 1** sort output on column 1 (default)
- 2** sort output on column 2
- 3** sort output on column 3

FILES

/tmp/crt??	temporaries
/usr/lib/aign	default assembler <i>ignore</i> file
/usr/lib/cign	default C <i>ignore</i> file
/usr/bin/crpost	post processor
/usr/bin/upost	post processor for -u option
/bin/sort	used to sort temporaries

SEE ALSO

as(I), cc(I)

BUGS

DATE(I)

11/1/73

DATE(I)

NAME

date – print and set the date

SYNOPSIS

date [mmddhhmm[yy]]

DESCRIPTION

If no argument is given, the current date is printed to the second. If an argument is given, the current date is set. The first *mm* is the month number; *dd* is the day number in the month; *hh* is the hour number (24 hour system); the second *mm* is the minute number; *yy* is the last 2 digits of the year number and is optional. For example:

date 10080045

sets the date to Oct 8, 12:45 AM. The current year is the default if no year is mentioned. The system operates in GMT. *Date* takes care of the conversion to and from local standard and daylight time.

DIAGNOSTICS

'bad conversion' if the argument is syntactically incorrect.

BUGS

NAME

db – debug

SYNOPSIS**db** [core [namelist]] [–]**DESCRIPTION**

Unlike many debugging packages (including DEC's ODT, on which *db* is loosely based), *db* is not loaded as part of the core image which it is used to examine; instead it examines files. Typically, the file will be either a core image produced after a fault or the binary output of the assembler. *Core* is the file being debugged; if omitted **core** is assumed. *Namelist* is a file containing a symbol table. If it is omitted, the symbol table is obtained from the file being debugged, or if not there from **a.out**. If no appropriate name list file can be found, *db* can still be used but some of its symbolic facilities become unavailable.

For the meaning of the optional third argument, see the last paragraph below.

The format for most *db* requests is an address followed by a one character command. Addresses are expressions built up as follows:

1. A name has the value assigned to it when the input file was assembled. It may be relocatable or not depending on the use of the name during the assembly.
2. An octal number is an absolute quantity with the appropriate value.
3. A decimal number immediately followed by '.' is an absolute quantity with the appropriate value.
4. An octal number immediately followed by **r** is a relocatable quantity with the appropriate value.
5. The symbol . indicates the current pointer of *db*. The current pointer is set by many *db* requests.
6. A * before an expression forms an expression whose value is the number in the word addressed by the first expression. A * alone is equivalent to '*.'
7. Expressions separated by + or blank are expressions with value equal to the sum of the components. At most one of the components may be relocatable.
8. Expressions separated by – form an expression with value equal to the difference to the components. If the right component is relocatable, the left component must be relocatable.
9. Expressions are evaluated left to right.

Names for registers are built in:

r0 ... r5
sp
pc
fr0 ... fr5

These may be examined. Their values are deduced from the contents of the stack in a core image file. They are meaningless in a file that is not a core image.

If no address is given for a command, the current address (also specified by “.”) is assumed. In general, “.” points to the last word or byte printed by *db*.

There are *db* commands for examining locations interpreted as numbers, machine instructions, ASCII characters, and addresses. For numbers and characters, either bytes or words may be examined. The following commands are used to examine the specified file.

- / The addressed word is printed in octal.
- \ The addressed byte is printed in octal.

- " The addressed word is printed as two ASCII characters.
- ˘ The addressed byte is printed as an ASCII character.
- ‘ The addressed word is printed in decimal.
- ? The addressed word is interpreted as a machine instruction and a symbolic form of the instruction, including symbolic addresses, is printed. Often, the result will appear exactly as it was written in the source program.
- & The addressed word is interpreted as a symbolic address and is printed as the name of the symbol whose value is closest to the addressed word, possibly followed by a signed offset.
- <nl>(i. e., the character “new line”) This command advances the current location counter “.” and prints the resulting location in the mode last specified by one of the above requests.
- ^ This character decrements “.” and prints the resulting location in the mode last selected one of the above requests. It is a converse to <nl>.
- % Exit.

Odd addresses to word-oriented commands are rounded down. The incrementing and decrementing of “.” done by the <nl> and ^ requests is by one or two depending on whether the last command was word or byte oriented.

The address portion of any of the above commands may be followed by a comma and then by an expression. In this case that number of sequential words or bytes specified by the expression is printed. “.” is advanced so that it points at the last thing printed.

There are two commands to interpret the value of expressions.

- = When preceded by an expression, the value of the expression is typed in octal. When not preceded by an expression, the value of “.” is indicated. This command does not change the value of “.”.
- : An attempt is made to print the given expression as a symbolic address. If the expression is relocatable, that symbol is found whose value is nearest that of the expression, and the symbol is typed, followed by a sign and the appropriate offset. If the value of the expression is absolute, a symbol with exactly the indicated value is sought and printed if found; if no matching symbol is discovered, the octal value of the expression is given.

The following command may be used to patch the file being debugged.

- ! This command must be preceded by an expression. The value of the expression is stored at the location addressed by the current value of “.”. The opcodes do not appear in the symbol table, so the user must assemble them by hand.

The following command is used after a fault has caused a core image file to be produced.

- \$ causes the fault type and the contents of the general registers and several other registers to be printed both in octal and symbolic format. The values are as they were at the time of the fault.

For some purposes, it is important to know how addresses typed by the user correspond with locations in the file being debugged. The mapping algorithm employed by *db* is non-trivial for two reasons: First, in an **a.out** file, there is a 20(8) byte header which will not appear when the file is loaded into core for execution. Therefore, apparent location 0 should correspond with actual file offset 20. Second, addresses in core images do not correspond with the addresses used by the program because in a core image there is a header containing the system’s per-process data for the dumped process, and also because the stack is stored contiguously with the text and data part of the core image rather than at the highest possible locations. *Db* obeys the following rules:

If exactly one argument is given, and if it appears to be an **a.out** file, the 20-byte header is skipped during addressing, i.e., 20 is added to all addresses typed. As a consequence, the header can be examined beginning at location -20.

If exactly one argument is given and if the file does not appear to be an **a.out** file, no mapping is done.

If zero or two arguments are given, the mapping appropriate to a core image file is employed. This means that locations above the program break and below the stack effectively do not exist (and are not, in fact, recorded in the core file). Locations above the user's stack pointer are mapped, in looking at the core file, to the place where they are really stored. The per-process data kept by the system, which is stored in the first part of the core file, cannot currently be examined (except by \$).

If one wants to examine a file which has an associated name list, but is not a core image file, the last argument “-” can be used (actually the only purpose of the last argument is to make the number of arguments not equal to two). This feature is used most frequently in examining the memory file /dev/mem.

SEE ALSO

as (I), core (V), a.out (V), od (I)

DIAGNOSTICS

“File not found” if the first argument cannot be read; otherwise “?”.

BUGS

There should be some way to examine the registers and other per-process data in a core image; also there should be some way of specifying double-precision addresses. It does not know yet about shared text segments.

NAME

dc – desk calculator

SYNOPSIS

dc [file]

DESCRIPTION

Dc is an arbitrary precision integer arithmetic package. The overall structure of *dc* is a stacking (reverse Polish) calculator. The following constructions are recognized by the calculator:

number The value of the number is pushed on the stack. A number is an unbroken string of the digits 0-9. It may be preceded by an underscore `_` to input a negative number.

`+`

`-`

`*`

`%`

`^`

The top two values on the stack are added (`+`), subtracted (`-`), multiplied (`*`), divided (`/`), remaindered (`%`), or exponentiated (`^`). The two entries are popped off the stack; the result is pushed on the stack in their place.

sx The top of the stack is popped and stored into a register named *x*, where *x* may be any character.

lx The value in register *x* is pushed on the stack. The register *x* is not altered. All registers start with zero value.

d The top value on the stack is pushed on the stack. Thus the top value is duplicated.

p The top value on the stack is printed. The top value remains unchanged.

f All values on the stack and in registers are printed.

q exits the program. If executing a string, the nesting level is popped by two.

x treats the top element of the stack as a character string and executes it as a string of dc commands.

[|...|] puts the bracketed ascii string onto the top of the stack.

`<x`

`=x`

`>x`

The top two elements of the stack are popped and compared. Register *x* is executed if they obey the stated relation.

v replaces the top element on the stack by its square root.

! interprets the rest of the line as a UNIX command.

c All values on the stack are popped.

i The top value on the stack is popped and used as the number radix for further input.

o The top value on the stack is popped and used as the number radix for further output.

z The stack level is pushed onto the stack.

? A line of input is taken from the input source (usually the console) and executed.

new-line ignored except as the name of a register or to end the response to a **?**.

space ignored except as the name of a register or to terminate a number.

If a file name is given, input is taken from that file until end-of-file, then input is taken from the console. An example which prints the first ten values of *n!* is

[la1+dsa*pla10>x]sx
0sa1
lxx

FILES

/etc/msh to implement '!'

DIAGNOSTICS

- (x) ? for unrecognized character x.
- (x) ? for not enough elements on the stack to do what was asked by command x.
- 'Out of space' when the free list is exhausted (too many digits).
- 'Out of headers' for too many numbers being kept around.
- 'Out of pushdown' for too many items on the stack.
- 'Nesting Depth' for too many levels of nested execution.

BUGS

NAME

dd – convert and copy a file

SYNOPSIS

dd [option=value] ...

DESCRIPTION

Dd copies the specified input file to the specified output with possible conversions. The input and output block size may be specified to take advantage of raw physical I/O.

<i>option</i>	<i>values</i>
if=	input file name; standard input is default
of=	output file name; standard output is default
ibs=	input block size (default 512)
obs=	output block size (default 512)
bs=	set both input and output block size, superseding <i>ibs</i> and <i>obs</i> ; also, if no conversion is specified, it is particularly efficient since no copy need be done
cbs= <i>n</i>	conversion buffer size
skip= <i>n</i>	skip <i>n</i> input records before starting copy
count= <i>n</i>	copy only <i>n</i> input records
conv=ascii	convert EBCDIC to ASCII
ebcdic	convert ASCII to EBCDIC
lcase	map alphabetic to lower case
ucase	map alphabetic to upper case
swab	swap every pair of bytes
noerror	do not stop processing on an error
sync	pad every input record to <i>ibs</i>
... , ...	several comma-separated conversions

Where sizes are specified, a number of bytes is expected. A number may end with **k**, **b** or **w** to specify multiplication by 1024, 512, or 2 respectively. Also a pair of numbers may be separated by **x** to indicate a product.

Cbs is used only if *ascii* or *ebcdic* conversion is specified. In the former case *cbs* characters are placed into the conversion buffer, converted to ASCII, and trailing blanks trimmed and new-line added before sending the line to the output. In the latter case ASCII characters are read into the conversion buffer, converted to EBCDIC, and blanks added to make up an output record of size *cbs*.

After completion, *dd* reports the number of whole and partial input and output blocks.

For example, to read an EBCDIC tape blocked ten 80-byte EBCDIC card images per record into the ASCII file *x*:

```
dd if=/dev/rmt0 of=x ibs=800 cbs=80 conv=ascii,lcase
```

Note the use of raw magtape. *Dd* is especially suited to I/O on the raw physical devices because it allows reading and writing in arbitrary record sizes.

SEE ALSO

cp (I)

BUGS

The ASCII/EBCDIC conversion tables are taken from the 256 character standard in the CACM Nov, 1968. It is not clear how this relates to real life.

NAME

diff – differential file comparator

SYNOPSIS

diff [-] name1 name2

DESCRIPTION

Diff tells what lines must be changed in two files to bring them into agreement. The normal output contains lines of these forms:

```
n1 a n3,n4
n1,n2 d n3
n1,n2 c n3,n4
```

These lines resemble *ed* commands to convert file *name1* into file *name2*. The numbers after the letters pertain to file *name2*. In fact, by exchanging ‘a’ for ‘d’ and reading backward one may ascertain equally how to convert file *name2* into *name1*. As in *ed*, identical pairs where $n1 = n2$ or $n3 = n4$ are abbreviated as a single number.

Following each of these lines come all the lines that are affected in the first file flagged by ‘*’, then all the lines that are affected in the second file flagged by ‘.’.

Under the – option, the output of *diff* is a script of *a*, *c* and *d* commands for the editor *ed*, which will change the contents of the first file into the contents of the second. In this connection, the following shell program may help maintain multiple versions of a file. Only an ancestral file (\$1) and a chain of version-to-version *ed* scripts (\$2,\$3,...) made by *diff* need be on hand. A ‘latest version’ appears on the standard output.

```
(cat $2 ... $9; echo "1,$p") | ed - $1
```

Diff does an optimal and unfailing job of detecting the file differences, and also reports these differences side-by-side. However, *diff* uses a quadratic algorithm that usually slows to a crawl on 2000-line files.

SEE ALSO

cmp (1), comm (1), ed (1)

DIAGNOSTICS

‘can’t open input’

‘arg count’

‘jackpot’ – To speed things up, the program uses hashing. You have stumbled on a case where there is a minuscule chance that this has resulted in an unnecessarily long list of differences being published. It’s a curio that we’d like to see.

BUGS

Editing scripts produced under the – option are naive about creating lines consisting of a single ‘.’.

NAME

dsw – delete interactively

SYNOPSIS

dsw [directory]

DESCRIPTION

For each file in the given directory (‘.’ if not specified) *dsw* types its name. If **y** is typed, the file is deleted; if **x**, *dsw* exits; if new-line, the file is not deleted; if anything else, *dsw* asks again.

SEE ALSO

rm(I)

BUGS

The name *dsw* is a carryover from the ancient past. Its etymology is amusing.

NAME

du - summarize disk usage

SYNOPSIS

du [**-s**] [**-a**] [name ...]

DESCRIPTION

Du gives the number of blocks contained in all files and (recursively) directories within each specified directory or file *name*. If *name* is missing, '.' is used.

The optional argument **-s** causes only the grand total to be given. The optional argument **-a** causes an entry to be generated for each file. Absence of either causes an entry to be generated for each directory only.

A file which has two links to it is only counted once.

BUGS

Non-directories given as arguments (not under **-a** option) are not listed.

Removable file systems do not work correctly since i-numbers may be repeated while the corresponding files are distinct. *Du* should maintain an i-number list per root directory encountered.

NAME

echo – echo arguments

SYNOPSIS

echo [arg ...]

DESCRIPTION

Echo writes all its arguments in order as a line on the standard output file. It is mainly useful for producing diagnostics in command files.

BUGS

NAME

ed – editor

SYNOPSIS

ed [-] [name]

DESCRIPTION

Ed is the standard text editor.

If a *name* argument is given, *ed* simulates an *e* command (see below) on the named file; that is to say, the file is read into *ed*'s buffer so that it can be edited. The optional - simulates an **os** command (see below) which suppresses the printing of characters counts by *e*, *r*, and *w* commands.

Ed operates on a copy of any file it is editing; changes made in the copy have no effect on the file until a *w* (write) command is given. The copy of the text being edited resides in a temporary file called the *buffer*. There is only one buffer.

Commands to *ed* have a simple and regular structure: zero or more *addresses* followed by a single character *command*, possibly followed by parameters to the command. These addresses specify one or more lines in the buffer. Every command which requires addresses has default addresses, so that the addresses can often be omitted.

In general, only one command may appear on a line. Certain commands allow the input of text. This text is placed in the appropriate place in the buffer. While *ed* is accepting text, it is said to be in *input mode*. In this mode, no commands are recognized; all input is merely collected. Input mode is left by typing a period '.' alone at the beginning of a line.

Ed supports a limited form of *regular expression* notation. A regular expression is an expression which specifies a set of strings of characters. A member of this set of strings is said to be *matched* by the regular expression. The regular expressions allowed by *ed* are constructed as follows:

1. An ordinary character (not one of those discussed below) is a regular expression and matches that character.
2. A circumflex '^' at the beginning of a regular expression matches the null character at the beginning of a line.
3. A currency symbol '\$' at the end of a regular expression matches the null character at the end of a line.
4. A period '.' matches any character but a new-line character.
5. A regular expression followed by an asterisk '*' matches any number of adjacent occurrences (including zero) of the regular expression it follows.
6. A string of characters enclosed in square brackets '[']' matches any character in the string but no others. If, however, the first character of the string is a circumflex '^' the regular expression matches any character but new-line and the characters in the string.
7. The concatenation of regular expressions is a regular expression which matches the concatenation of the strings matched by the components of the regular expression.
8. The null regular expression standing alone is equivalent to the last regular expression encountered.

Regular expressions are used in addresses to specify lines and in one command (see *s* below) to specify a portion of a line which is to be replaced.

If it is desired to use one of the regular expression metacharacters as an ordinary character, that character may be preceded by '\'. This also applies to the character bounding the regular expression (often '/') and to '\' itself.

Addresses are constructed as follows. To understand addressing in *ed* it is necessary to know that at any time there is a *current line*. Generally speaking, the current line is the last line affected by a command; however, the exact effect on the current line by each command is discussed under the description of the command.

1. The character ‘.’ addresses the current line.
2. The character ‘^’ addresses the line immediately before the current line.
3. The character ‘\$’ addresses the last line of the buffer.
4. A decimal number *n* addresses the *n*-th line of the buffer.
5. ‘*x*’ addresses the line associated (marked) with the mark name character *x* which must be a printable character. Lines are marked with the *k* command described below.
6. A regular expression enclosed in slashes ‘/’ addresses the first line found by searching toward the end of the buffer and stopping at the first line containing a string matching the regular expression. If necessary the search wraps around to the beginning of the buffer.
7. A regular expression enclosed in queries ‘?’ addresses the first line found by searching toward the beginning of the buffer and stopping at the first line found containing a string matching the regular expression. If necessary the search wraps around to the end of the buffer.
8. An address followed by a plus sign ‘+’ or a minus sign ‘-’ followed by a decimal number specifies that address plus (resp. minus) the indicated number of lines. The plus sign may be omitted.

Commands may require zero, one, or two addresses. Commands which require no addresses regard the presence of an address as an error. Commands which accept one or two addresses assume default addresses when insufficient are given. If more addresses are given than such a command requires, the last one or two (depending on what is accepted) are used.

Addresses are separated from each other typically by a comma ‘,’. They may also be separated by a semicolon ‘;’. In this case the current line ‘.’ is set to the previous address before the next address is interpreted. This feature can be used to determine the starting line for forward and backward searches (‘/’, ‘?’). The second address of any two-address sequence must correspond to a line following the line corresponding to the first address.

In the following list of *ed* commands, the default addresses are shown in parentheses. The parentheses are not part of the address, but are used to show that the given addresses are the default.

As mentioned, it is generally illegal for more than one command to appear on a line. However, any command may be suffixed by ‘p’ (for ‘print’). In that case, the current line is printed after the command is complete.

(.)a
<text>

•

The append command reads the given text and appends it after the addressed line. ‘.’ is left on the last line input, if there were any, otherwise at the addressed line. Address ‘0’ is legal for this command; text is placed at the beginning of the buffer.

(. , .)c
<text>

•

The change command deletes the addressed lines, then accepts input text which replaces these lines. ‘.’ is left at the last line input; if there were none, it is left at the first line not changed.

(. , .)d

The delete command deletes the addressed lines from the buffer. The line originally after the last line deleted becomes the current line; if the lines deleted were originally at the end, the new last line becomes the current line.

e filename

The edit command causes the entire contents of the buffer to be deleted, and then the named file to be read in. '.' is set to the last line of the buffer. The number of characters read is typed. 'filename' is remembered for possible use as a default file name in a subsequent *r* or *w* command.

f filename

The filename command prints the currently remembered file name. If 'filename' is given, the currently remembered file name is changed to 'filename'.

(1,\$)g/regular expression/command list

In the global command, the first step is to mark every line which matches the given regular expression. Then for every such line, the given command list is executed with '.' initially set to that line. A single command or the first of multiple commands appears on the same line with the global command. All lines of a multi-line list except the last line must be ended with '\'. *A*, *i*, and *c* commands and associated input are permitted; the '.' terminating input mode may be omitted if it would be on the last line of the command list. The (global) commands, *g*, and *v*, are not permitted in the command list.

**(.)i
<text>**

This command inserts the given text before the addressed line. '.' is left at the last line input; if there were none, at the addressed line. This command differs from the *a* command only in the placement of the text.

(.)kx

The mark command associates or marks the addressed line with the single character mark name *x*. The ten most recent mark names are remembered. The current mark names may be printed with the *n* command.

(...)l

The list command will print the addressed lines in a way that is unambiguous: Non-graphic characters are printed in octal, prefixed characters are overstruck with a circumflex, and long lines are folded.

(...)ma

The move command will reposition the addressed lines after the line addressed by *a*. The last of the moved lines becomes the current line.

n

The *n* command will print the current mark names.

os**ov**

After *os* character counts printed by *e*, *r*, and *w* are suppressed. *ov* turns them back on.

(...)p

The print command prints the addressed lines. '.' is left at the last line printed. The *p* command may be placed on the same line after any command.

q

The quit command causes *ed* to exit. No automatic write of a file is done.

(\$)r filename

The read command reads in the given file after the addressed line. If no file name is given, the remembered file name, if any, is used (see *e* and *f* commands). The remembered file name is not changed unless 'filename' is the very first file name mentioned. Address '0' is legal for *r* and causes the file to be read at the beginning of the buffer. If the read is successful, the number of characters read is typed. '.' is left at the last line read in from the file.

- (. . .) s/regular expression/replacement/ or,
- (. . .) s/regular expression/replacement/g

The substitute command searches each addressed line for an occurrence of the specified regular expression. On each line in which a match is found, all matched strings are replaced by the replacement specified, if the global replacement indicator 'g' appears after the command. If the global indicator does not appear, only the first occurrence of the matched string is replaced. It is an error for the substitution to fail on all addressed lines. Any character other than space or new-line may be used instead of '/' to delimit the regular expression and the replacement. '.' is left at the last line substituted.

An ampersand '&' appearing in the replacement is replaced by the regular expression that was matched. The special meaning of '&' in this context may be suppressed by preceding it by '\'.

- (1,\$) v/regular expression/command list

This command is the same as the global command except that the command list is executed with '.' initially set to every line *except* those matching the regular expression.

- (1,\$) w filename

The write command writes the addressed lines onto the given file. If the file does not exist, it is created mode 666 (readable and writeable by everyone). The remembered file name is *not* changed unless 'filename' is the very first file name mentioned. If no file name is given, the remembered file name, if any, is used (see *e* and *f* commands). '.' is unchanged. If the command is successful, the number of characters written is typed.

- (\$)=

The line number of the addressed line is typed. '.' is unchanged by this command.

- !UNIX command

The remainder of the line after the '!' is sent to UNIX to be interpreted as a command. '.' is unchanged. The entire shell syntax is not recognized. See msh(VII) for the restrictions.

- (.+1) <newline>

An address alone on a line causes the addressed line to be printed. A blank line alone is equivalent to '.+1p'; it is useful for stepping through text.

If an interrupt signal (ASCII DEL) is sent, *ed* will print a '?' and return to its command level.

If invoked with the command name '-', (see init(VII)) *ed* will sign on with the message 'Editing system' and print '*' as the command level prompt character.

Ed has size limitations on the maximum number of lines that can be edited, on the maximum number of characters in a line, in a global's command list, in a remembered file name, and in the size of the temporary file. The current sizes are: 4000 lines per file, 512 characters per line, 256 characters per global command list, 64 characters per file name, and 64K characters in the temporary file (see BUGS).

FILES

/tmp/etm?, temporary
 /etc/msh, to implement the '!' command.

DIAGNOSTICS

'?' for errors in commands; 'TMP' for temporary file overflow.

SEE ALSO

A Tutorial Introduction to the ED Text Editor (internal memorandum)

BUGS

The temporary file can grow to no more than 64K bytes.

NAME

eqn — typeset mathematics

SYNOPSIS

eqn [file] ...

DESCRIPTION

Eqn is a troff (I) preprocessor for typesetting mathematics on the Graphics Systems phototypesetter. Usage is almost always

eqn file ... | troff

If no files are specified, *eqn* reads from the standard input. A line beginning with “.EQ” marks the start of an equation; the end of an equation is marked by a line beginning with “.EN”. Neither of these lines is altered or defined by *eqn*, so you can define them yourself to get centering, numbering, etc. All other lines are treated as comments, and passed through untouched.

Spaces, tabs, newlines, braces, double quotes, tilde and circumflex are the only delimiters. Braces “{ }” are used for grouping. Use tildes “~” to get extra spaces in an equation.

Subscripts and superscripts are produced with the keywords **sub** and **sup**. Thus *x sub i* makes x_i , *a sub i sup 2* produces a_i^2 , and *e sup {x sup 2 + y sup 2}* gives $e^{x^2+y^2}$. Fractions are made with **over**. *a over b* is $\frac{a}{b}$ and *1 over sqrt {ax sup 2 + bx + c}* is $\frac{1}{\sqrt{ax^2+bx+c}}$. **sqrt** makes square roots.

The keywords **from** and **to** introduce lower and upper limits on arbitrary things: $\lim_{n \rightarrow \infty} \sum_0^n x_i$ is made with *lim from {n-> inf} sum from 0 to n x sub i*. Left and right brackets, braces, etc., of the right height are made with **left** and **right**: *left [x sup 2 + y sup 2 over alpha right] ~ = 1* produces $\left[x^2 + \frac{y^2}{\alpha} \right] = 1$. The **right** clause is optional.

Vertical piles of things are made with **pile**, **lpile**, **cpile**, and **rpile**: *pile {a above b above c}* produces $\begin{matrix} a \\ b \\ c \end{matrix}$. There can be an arbitrary number of elements in a pile. **lpile** left-justifies, **pile** and **cpile** center, with different vertical spacing, and **rpile** right justifies.

Diacritical marks are made with **dot**, **dotdot**, **hat**, **bar**: *x dot = f(t)* *bar* is $\dot{x} = \overline{f(t)}$. Default sizes and fonts can be changed with **size n** and various of **roman**, **italic**, and **bold**.

Keywords like *sum* (\sum) *int* (\int) *inf* (∞) and shorthands like \geq (\geq) \rightarrow (\rightarrow), \neq (\neq), are recognized. Spell out Greek letters in the desired case, as in *alpha*, *GAMMA*. Mathematical words like sin, cos, log are made Roman automatically. Troff (I) four-character escapes like \backslash (bs \backslash) can be used anywhere. Strings enclosed in double quotes “...” are passed through untouched.

SEE ALSO

A System for Typesetting Mathematics (Computer Science Technical Report #17, Bell Laboratories, 1974.)
TROFF Users’ Manual (internal memorandum)
TROFF Made Trivial (internal memorandum)
troff (I), neqn (I)

BUGS

Undoubtedly. Watch out for small or large point sizes — it’s tuned too well for size 10. Be cautious if inserting horizontal or vertical motions, and of backslashes in general.

NAME

exit – terminate command file

SYNOPSIS

exit

DESCRIPTION

Exit performs a **seek** to the end of its standard input file. Thus, if it is invoked inside a file of commands, upon return from **exit** the shell will discover an end-of-file and terminate.

SEE ALSO

if (I), goto (I), sh (I)

BUGS

NAME

fc – fortran compiler

SYNOPSIS

fc [**-c**] sfile1.f ... ofile1 ...

DESCRIPTION

Fc is the UNIX Fortran compiler. It accepts three types of arguments:

Arguments whose names end with ‘.f’ are assumed to be Fortran source program units; they are compiled, and the object program is left on the file sfile1.o (i.e. the file whose name is that of the source with ‘.o’ substituted for ‘.f’).

Other arguments (except for **-c**) are assumed to be either loader flags, or object programs, typically produced by an earlier *fc* run, or perhaps libraries of Fortran-compatible routines. These programs, together with the results of any compilations specified, are loaded (in the order given) to produce an executable program with name **a.out**.

The **-c** argument suppresses the loading phase, as does any syntax error in any of the routines being compiled.

The following is a list of differences between *fc* and ANSI standard Fortran (also see the BUGS section):

1. Arbitrary combination of types is allowed in expressions. Not all combinations are expected to be supported at runtime. All of the normal conversions involving integer, real, double precision and complex are allowed.
2. DEC’s **implicit** statement is recognized. E.g.: **implicit integer /i-n/**
3. The types doublecomplex, logical*1, integer*1, integer*2 and real*8 (double precision) are supported.
4. **&** as the first character of a line signals a continuation card.
5. **c** as the first character of a line signals a comment.
6. All keywords are recognized in lower case.
7. The notion of ‘column 7’ is not implemented.
8. G-format input is free form– leading blanks are ignored, the first blank after the start of the number terminates the field.
9. A comma in any numeric or logical input field terminates the field.
10. There is no carriage control on output.
11. A sequence of *n* characters in double quotes “” is equivalent to *n* **h** followed by those characters.
12. In **data** statements, a hollerith string may initialize an array or a sequence of array elements.
13. The number of storage units requested by a binary **read** must be identical to the number contained in the record being read.
14. If the first character in an input file is ‘#’, a preprocessor identical to the C preprocessor is called, which implements “#define” and “#include” preprocessor statements. (See the C reference manual for details.) The preprocessor does not recognize Hollerith strings written with *n* **h**.

In I/O statements, only unit numbers 0-19 are supported. Unit number *n* refers to file *fortn*; (e.g. unit 9 is file ‘fort09’). For input, the file must exist; for output, it will be created. Unit 5 is permanently associated with the standard input file; unit 6 with the standard output file. Also see *setfil* (III) for a way to associate unit numbers with named files.

FILES

file.f	input file
a.out	loaded output
f.tmp[123]	temporary (deleted)
/usr/fort/fc1	compiler proper
/lib/fr0.o	runtime startoff
/lib/filib.a	interpreter library
/lib/libf.a	builtin functions, etc.
/lib/liba.a	system library

SEE ALSO

ANSI standard, ld (I) for loader flags

Also see the writeups on the precious few non-standard Fortran subroutines, *ierror* and *setfil* (III)

DIAGNOSTICS

Compile-time diagnostics are given in English, accompanied if possible with the offending line number and source line with an underscore where the error occurred. Runtime diagnostics are given by number as follows:

- 1 invalid log argument
- 2 bad arg count to amod
- 3 bad arg count to atan2
- 4 excessive argument to cabs
- 5 exp too large in cexp
- 6 bad arg count to cmplx
- 7 bad arg count to dim
- 8 excessive argument to exp
- 9 bad arg count to idim
- 10 bad arg count to isign
- 11 bad arg count to mod
- 12 bad arg count to sign
- 13 illegal argument to sqrt
- 14 assigned/computed goto out of range
- 15 subscript out of range
- 16 real**real overflow
- 17 (negative real)**real
- 100 illegal I/O unit number
- 101 inconsistent use of I/O unit
- 102 cannot create output file
- 103 cannot open input file
- 104 EOF on input file
- 105 illegal character in format
- 106 format does not begin with (
- 107 no conversion in format but non-empty list
- 108 excessive parenthesis depth in format
- 109 illegal format specification
- 110 illegal character in input field
- 111 end of format in hollerith specification
- 112 bad argument to setfil
- 120 bad argument to ierror
- 999 unimplemented input conversion

Any of these errors can be caught by the program; see *ierror* (III).

BUGS

The following is a list of those features not yet implemented:

arithmetic statement functions

scale factors on input

Backspace statement.

NAME

fed – edit associative memory for form letter

SYNOPSIS

fed

DESCRIPTION

Fed is used to edit a form letter associative memory file, **form.m**, which consists of named strings. Commands consist of single letters followed by a list of string names separated by a single space and ending with a new line. The conventions of the Shell with respect to ‘*’ and ‘?’ hold for all commands but **m**. The commands are:

e name ...

Fed writes the string whose name is *name* onto a temporary file and executes *ed*. On exit from the *ed* the temporary file is copied back into the associative memory. Each argument is operated on separately. Be sure to give an *ed w* command (without a filename) to rewrite *fed*’s temporary file before quitting out of *ed*.

d [name ...]

deletes a string and its name from the memory. When called with no arguments **d** operates in a verbose mode typing each string name and deleting only if a **y** is typed. A **q** response returns to *fed*’s command level. Any other response does nothing.

m name1 name2 ...

(move) changes the name of name1 to name2 and removes previous string name2 if one exists. Several pairs of arguments may be given. Literal strings are expected for the names.

n [name ...]

(names) lists the string names in the memory. If called with the optional arguments, it just lists those requested.

p name ...

prints the contents of the strings with names given by the arguments.

q

returns to the system.

c [**p**] [**f**]

checks the associative memory file for consistency and reports the number of free headers and blocks. The optional arguments do the following:

p causes any unaccounted-for string to be printed.

f fixes broken memories by adding unaccounted-for headers to free storage and removing references to released headers from associative memory.

FILES

/tmp/ftmp?	temporary
form.m	associative memory

SEE ALSO

form(I), ed(I), sh(I)

WARNING

It is legal but unwise to have string names with blanks, ‘.’ or ‘?’ in them.

BUGS

NAME

find – find files

SYNOPSIS**find** pathname expression**DESCRIPTION**

Find recursively descends the directory hierarchy from *pathname* seeking files that match a boolean *expression* written in the primaries given below. In the descriptions, the argument *n* is used as a decimal integer where *+n* means more than *n*, *-n* means less than *n* and *n* means exactly *n*.

- name** filename True if the *filename* argument matches the current file name. Normal *Shell* argument syntax may be used if escaped (watch out for '[', '?' and '*').
- perm** onum True if the file permission flags exactly match the octal number *onum* (see *chmod*(I)). If *onum* is prefixed by a minus sign, more flag bits (017777, see *stat*(II)) become significant and the flags are compared: *(flags&onum)==onum*.
- type** *c* True if the type of the file is *c*, where *c* is **b**, **c**, **d** or **f** for block special file, character special file, directory or plain file.
- links** *n* True if the file has *n* links.
- user** *uname* True if the file belongs to the user *uname*.
- group** *gname* As it is for **-user** so shall it be for **-group** (someday).
- size** *n* True if the file is *n* blocks long (512 bytes per block).
- atime** *n* True if the file has been accessed in *n* days.
- mtime** *n* True if the file has been modified in *n* days.
- exec** command True if the executed command returns exit status zero (most commands do). The end of the command is punctuated by an escaped semicolon. A command argument '{ }' is replaced by the current pathname.
- ok** command Like **-exec** except that the generated command line is printed with a question mark first, and is executed only if the user responds **y**.
- print** Always true; causes the current pathname to be printed.

The primaries may be combined with these operators (ordered by precedence):

- !** prefix *not*
- a** infix *and*, second operand evaluated only if first is true
- o** infix *or*, second operand evaluated only if first is false
- (expression) parentheses for grouping. (Must be escaped.)

To remove files named 'a.out' and '*.o' not accessed for a week:

```
find / "(" -name a.out -o -name "*.o" ")" -a -atime +7 -a -exec rm { } ";"
```

FILES

/etc/passwd

SEE ALSO

sh (I), if(I), file system (V)

BUGS

There is no way to check device type.
Syntax should be reconciled with *if*.

NAME

form – form letter generator

SYNOPSIS

form proto arg ...

DESCRIPTION

Form generates a form letter from a prototype letter, an associative memory, arguments and in a special case, the current date.

If *form* is invoked with the *proto* argument *x*, the associative memory is searched for an entry with name *x* and the contents filed under that name are used as the prototype. If the search fails, the message '[x:]' is typed on the console and whatever text is typed in from the console, terminated by two new lines, is used as the prototype. If the prototype argument is missing, '{letter}' is assumed.

Basically, *form* is a copy process from the prototype to the output file. If an element of the form [*n*] (where *n* is a digit from 1 to 9) is encountered, the *n*-th argument is inserted in its place, and that argument is then rescanned. If [0] is encountered, the current date is inserted. If the desired argument has not been given, a message of the form '[*n*:]' is typed. The response typed in then is used for that argument.

If an element of the form [*name*] or {*name*} is encountered, the *name* is looked up in the associative memory. If it is found, the contents of the memory under this *name* replaces the original element (again rescanned). If the *name* is not found, a message of the form '[*name*:]' is typed. The response typed in is used for that element. The response is entered in the memory under the name if the name is enclosed in []. The response is not entered in the memory but is remembered for the duration of the letter if the name is enclosed in { }.

In both of the above cases, the response is typed in by entering arbitrary text terminated by two new lines. Only the first of the two new lines is passed with the text.

If one of the special characters [{}]\ is preceded by a \, it loses its special character.

If a file named 'forma' already exists in the user's directory, 'formb' is used as the output file and so forth to 'formz'.

The file 'form.m' is created if none exists. Because form.m is operated on by the disc allocator, it should only be changed by using *fed*, the form letter editor, or *form*.

FILES

form.m associative memory
form? output file (read only)

SEE ALSO

fed (I), roff (I)

BUGS

An unbalanced] or } acts as an end of file but may add a few strange entries to the associative memory.

NAME

`goto` — command transfer

SYNOPSIS

`goto` *label*

DESCRIPTION

Goto is only allowed when the Shell is taking commands from a file. The file is searched from the beginning for a line beginning with `:` followed by one or more spaces followed by the *label*. If such a line is found, the *goto* command returns. Since the read pointer in the command file points to the line after the label, the effect is to cause the Shell to transfer to the labelled line.

`:` is a do-nothing command that is ignored by the Shell and only serves to place a label.

SEE ALSO

`sh(I)`

BUGS

NAME

grep – search a file for a pattern

SYNOPSIS

grep [**-v**] [**-b**] [**-c**] [**-n**] *expression* [*file*] ...

DESCRIPTION

Grep will search the input files (standard input default) for each line containing the regular expression. Normally, each line found is printed on the standard output. If the **-v** flag is used, all lines but those matching are printed. If the **-c** flag is used, each line printed is preceded by its line number. If the **-n** flag is used, each line is preceded by the name of the input file and its relative line number in that file. If the **-b** flag is used, each line is preceded by the block number on which it was found. This is sometimes useful in locating disk block numbers by context. If interrupt is hit, the file and line number last searched is printed.

For a complete description of the regular expression, see *ed* (I). Care should be taken when using the characters \$ * [^ | () and \ in the regular expression as they are also meaningful to the Shell. It is generally necessary to enclose the entire *expression* argument in quotes.

SEE ALSO

ed (I), *sh* (I)

BUGS

Lines are limited to 256 characters; longer lines are truncated.

NAME

if – conditional command

SYNOPSIS

if *expr* *command* [*arg ...*]

DESCRIPTION

If evaluates the expression *expr*, and if its value is true, executes the given *command* with the given arguments.

The following primitives are used to construct the *expr*:

-r *file* true if the file exists and is readable.

-w *file* true if the file exists and is writable.

s1 = *s2* true if the strings *s1* and *s2* are equal.

s1 != *s2* true if the strings *s1* and *s2* are not equal.

{ *command* } The bracketed command is executed to obtain the exit status. Status zero is considered *true*. The command must not be another *if*.

These primaries may be combined with the following operators:

! unary negation operator

-a binary *and* operator

-o binary *or* operator

(*expr*) parentheses for grouping.

-a has higher precedence than **-o**. Notice that all the operators and flags are separate arguments to *if* and hence must be surrounded by spaces. Notice also that parentheses are meaningful to the Shell and must be escaped.

SEE ALSO

sh (I), find (I)

BUGS

NAME

kill – do in an unwanted process

SYNOPSIS

kill [-signo] processid ...

DESCRIPTION

Kills the specified processes. The processid of each asynchronous process started with ‘&’ is reported by the shell. Processid’s can also be found by using *ps* (I).

The killed process must have been started from the same typewriter as the current user, unless he is the super-user.

If a signal number preceded by “-” is given as first argument, that signal is sent instead of *kill* (see *signal* (II)).

SEE ALSO

ps (I), *sh* (I), *signal* (II)

BUGS

Clearly people should only be allowed to kill processes owned by them, and having the same typewriter is neither necessary nor sufficient.

NAME

ld - link editor

SYNOPSIS

ld [**-sulxrnd**] name ...

DESCRIPTION

Ld combines several object programs into one; resolves external references; and searches libraries. In the simplest case the names of several object programs are given, and *ld* combines them, producing an object module which can be either executed or become the input for a further *ld* run. (In the latter case, the **-r** option must be given to preserve the relocation bits.) The output of *ld* is left on **a.out**. This file is executable only if no errors occurred during the load.

The argument routines are concatenated in the order specified. The entry point of the output is the beginning of the first routine.

If any argument is a library, it is searched exactly once at the point it is encountered in the argument list. Only those routines defining an unresolved external reference are loaded. If a routine from a library references another routine in the library, the referenced routine must appear after the referencing routine in the library. Thus the order of programs within libraries is important.

Ld understands several flag arguments which are written preceded by a '-'. Except for **-l**, they should appear before the file names.

- s** 'squash' the output, that is, remove the symbol table and relocation bits to save space (but impair the usefulness of the debugger). This information can also be removed by *strip*.
- u** take the following argument as a symbol and enter it as undefined in the symbol table. This is useful for loading wholly from a library, since initially the symbol table is empty and an unresolved reference is needed to force the loading of the first routine.
- l** This option is an abbreviation for a library name. **-l** alone stands for '/lib/liba.a', which is the standard system library for assembly language programs. **-lx** stands for '/lib/libx.a' where *x* is any character. A library is searched when its name is encountered, so the placement of a **-l** is significant.
- x** do not preserve local (non-globl) symbols in the output symbol table; only enter external symbols. This option saves some space in the output file.
- r** generate relocation bits in the output file so that it can be the subject of another *ld* run. This flag also prevents final definitions from being given to common symbols.
- d** force definition of common storage even if the **-r** flag is present (used for reloc (VIII)).
- n** Arrange that when the output file is executed, the text portion will be read-only and shared among all users executing the file. This involves moving the data areas up the the first possible 4K word boundary following the end of the text.

FILES

/lib/lib?.a libraries
a.out output file

SEE ALSO

as(I), ar(I)

BUGS

NAME

ln — make a link

SYNOPSIS

ln name1 [name2]

DESCRIPTION

A link is a directory entry referring to a file; the same file (together with its size, all its protection information, etc) may have several links to it. There is no way to distinguish a link to a file from its original directory entry; any changes in the file are effective independently of the name by which the file is known.

Ln creates a link to an existing file *name1*. If *name2* is given, the link has that name; otherwise it is placed in the current directory and its name is the last component of *name1*.

It is forbidden to link to a directory or to link across file systems.

SEE ALSO

rm(I)

BUGS

There is nothing particularly wrong with *ln*, but *tp* doesn't understand about links and makes one copy for each name by which a file is known; thus if the tape is extracted several copies are re-stored and the information that links were involved is lost.

NAME

login – sign onto UNIX

SYNOPSIS

login [username]

DESCRIPTION

The *login* command is used when a user initially signs onto UNIX, or it may be used at any time to change from one user to another. The latter case is the one summarized above and described here. See ‘How to Get Started’ for how to dial up initially.

If *login* is invoked without an argument, it will ask for a user name, and, if appropriate, a password. Echoing is turned off (if possible) during the typing of the password, so it will not appear on the written record of the session.

After a successful login, accounting files are updated and the user is informed of the existence of *mailbox* and message-of-the-day files.

Login is recognized by the Shell and executed directly (without forking).

FILES

/tmp/utmp	accounting
/usr/adm/wtmp	accounting
mailbox	mail
/etc/motd	message-of-the-day
/etc/passwd	password file

SEE ALSO

init (VIII), getty (VIII), mail (I)

DIAGNOSTICS

‘login incorrect,’ if the name or the password is bad. ‘No Shell,’ ‘cannot open password file,’ ‘no directory’: consult a UNIX programming counselor.

BUGS

If the first login is unsuccessful, it tends to go into a state where it won’t accept a correct login. Hit EOT and try again.

NAME

`lpr` – on line print

SYNOPSIS

`lpr` [-] [+] [+-]file ...

DESCRIPTION

Lpr arranges to have the line printer daemon print the file arguments.

Normally, each file is printed in the state it is found when the line printer daemon reads it. If a particular file argument is preceded by +, or a preceding argument of + has been encountered, then *lpr* will make a copy for the daemon to print. If the file argument is preceded by -, or a preceding argument of - has been encountered, then *lpr* will unlink (remove) the file.

If there are no arguments, then the standard input is read and on-line printed. Thus *lpr* may be used as a filter.

FILES

<code>/usr/lpd/*</code>	spool area
<code>/etc/passwd</code>	personal ident cards
<code>/etc/lpd</code>	daemon

SEE ALSO

`lpd` (VIII), `passwd` (V)

BUGS

NAME

ls – list contents of directory

SYNOPSIS

ls [**-ltasdrui**] name ...

DESCRIPTION

For each directory argument, *ls* lists the contents of the directory; for each file argument, *ls* repeats its name and any other information requested. The output is sorted alphabetically by default. When no argument is given, the current directory is listed. When several arguments are given, the arguments are first sorted appropriately, but file arguments appear before directories and their contents. There are several options:

- l** list in long format, giving mode, number of links, owner, size in bytes, and time of last modification for each file. (See below.) If the file is a special file the size field will instead contain the major and minor device numbers.
- t** sort by time modified (latest first) instead of by name, as is normal
- a** list all entries; usually those beginning with '.' are suppressed
- s** give size in blocks for each entry
- d** if argument is a directory, list only its name, not its contents (mostly used with **-l** to get status on directory)
- r** reverse the order of sort to get reverse alphabetic or oldest first as appropriate
- u** use time of last access instead of last modification for sorting (**-t**) or printing (**-l**)
- i** print i-number in first column of the report for each file listed
- f** force each argument to be interpreted as a directory and list the name found in each slot. This option turns off **-l**, **-t**, **-s**, and **-r**, and turns on **-a**; the order is the order in which entries appear in the directory.

The mode printed under the **-l** option contains 10 characters which are interpreted as follows: the first character is

- d** if the entry is a directory;
- b** if the entry is a block-type special file;
- c** if the entry is a character-type special file;
- if the entry is a plain file.

The next 9 characters are interpreted as three sets of three bits each. The first set refers to owner permissions; the next to permissions to others in the same user-group; and the last to all others. Within each set the three characters indicate permission respectively to read, to write, or to execute the file as a program. For a directory, 'execute' permission is interpreted to mean permission to search the directory for a specified file. The permissions are indicated as follows:

- r** if the file is readable
- w** if the file is writable
- x** if the file is executable
- if the indicated permission is not granted

Finally, the group-execute permission character is given as **s** if the file has set-group-ID mode; likewise the user-execute permission character is given as **s** if the file has set-user-ID mode.

FILES

/etc/passwd to get user ID's for **ls -l**.

BUGS

NAME

mail – send mail to another user

SYNOPSIS

mail [**-yn**]

mail letter person ...

mail person

DESCRIPTION

Mail without an argument searches for a file called *mailbox*, prints it if present, and asks if it should be saved. If the answer is **y**, the mail is renamed *mbox*, otherwise it is deleted. *Mail* with a **-y** or **-n** argument works the same way, except that the answer to the question is supplied by the argument.

When followed by the names of a letter and one or more people, the letter is prepended to each person's *mailbox*. When a *person* is specified without a *letter*, the letter is taken from the sender's standard input up to an EOT. Each letter is preceded by the sender's name and a post-mark.

A *person* is either a user name recognized by **login**, in which case the mail is sent to the default working directory of that user, or the path name of a directory, in which case *mailbox* in that directory is used.

When a user logs in he is informed of the presence of mail.

FILES

/etc/passwd	to identify sender and locate persons
mailbox	input mail
mbox	saved mail
/tmp/mtm?	temp file

SEE ALSO

login (I)

BUGS

NAME

man – run off section of UNIX manual

SYNOPSIS

man [section] [title ...]

DESCRIPTION

Man is a shell command file that will locate and run off one or more sections of this manual. *Section* is the section number of the manual, as an Arabic not Roman numeral, and is optional. *Title* is one or more section names; these names bear a generally simple relation to the page captions in the manual. If the *section* is missing, **1** is assumed. For example,

man man

would reproduce this page.

FILES

/usr/man/man?/*

BUGS

The manual is supposed to be reproducible either on the phototypesetter or on a typewriter. However, on a typewriter some information is necessarily lost.

NAME

mesg – permit or deny messages

SYNOPSIS

mesg [**n**] [**y**]

DESCRIPTION

Mesg with argument **n** forbids messages via *write* by revoking non-user write permission on the user's typewriter. *Mesg* with argument **y** reinstates permission. All by itself, *mesg* reverses the current permission. In all cases the previous state is reported.

FILES

/dev/tty?

SEE ALSO

write (I)

DIAGNOSTICS

'?' if the standard input file is not a typewriter

BUGS

NAME

mkdir — make a directory

SYNOPSIS

mkdir dirname ...

DESCRIPTION

Mkdir creates specified directories in mode 777. The standard entries ‘.’ and ‘..’ are made automatically.

SEE ALSO

rmdir(I)

BUGS

NAME

mv — move or rename a file

SYNOPSIS

mv name1 name2

DESCRIPTION

Mv changes the name of *name1* to *name2*. If *name2* is a directory, *name1* is moved to that directory with its original file-name. Directories may only be moved within the same parent directory (just renamed).

If *name2* already exists, it is removed before *name1* is renamed. If *name2* has a mode which forbids writing, *mv* prints the mode and reads the standard input to obtain a line; if the line begins with **y**, the move takes place; if not, *mv* exits.

If *name2* would lie on a different file system, so that a simple rename is impossible, *mv* copies the file and deletes the original.

BUGS

It should take a **-f** flag, like *rm*, to suppress the question if the target exists and is not writable.

NAME

`neqn` – typeset mathematics on terminal

SYNOPSIS

`neqn` [file] ...

DESCRIPTION

Neqn is an *nroff* (I) preprocessor. The input language is the same as that of *eqn* (I). Normal usage is almost always

`neqn file ... | nroff`

Output is meant for terminals with forward and reverse capabilities, such as the Model 37 teletype or GSI terminal.

If no arguments are specified, *neqn* reads the standard input, so it may be used as a filter.

SEE ALSO

eqn (I), *gsi* (VI)

BUGS

Because of some interactions with *nroff* there may not always be enough space left before and after lines containing equations.

NAME

nice – run a command at low priority

SYNOPSIS

nice command [arguments]

DESCRIPTION

Nice executes *command* at low priority.

SEE ALSO

nohup(I), nice(II)

BUGS

NAME

nm - print name list

SYNOPSIS

nm [**-cnru**] [name]

DESCRIPTION

Nm prints the symbol table from the output file of an assembler or loader run. Each symbol name is preceded by its value (blanks if undefined) and one of the letters **U** (undefined) **A** (absolute) **T** (text segment symbol), **D** (data segment symbol), **B** (bss segment symbol), or **C** (common symbol). If the symbol is local (non-external) the type letter is in lower case. The output is sorted alphabetically.

If no file is given, the symbols in **a.out** are listed.

Options are:

- c** list only C-style external symbols, that is those beginning with underscore ‘_’.
- n** sort by value instead of by name
- r** sort in reverse order
- u** print only undefined symbols.

FILES

a.out

BUGS

NAME

nohup – run a command immune to hangups

SYNOPSIS

nohup command [arguments]

DESCRIPTION

Nohup executes *command* with hangups, quits and interrupts all ignored.

SEE ALSO

nice(I), signal(II)

BUGS

NAME

nroff - format text

SYNOPSIS

nroff [*+n*] [*-n*] [*-nn*] [*-mx*] [*-s*] [*-h*] [*-q*] [*-i*] files

DESCRIPTION

Nroff formats text according to control lines embedded in the text files. *Nroff* will read the standard input if no file arguments are given. The non-file option arguments are interpreted as follows:

- +n* Output will commence at the first page whose page number is *n* or larger
- n* will cause printing to stop after page *n*.
- nn* First generated (not necessarily printed) page is given number *n*; simulates “.pn *n*”.
- mx* Prepends a standard macro file; simulates “.so /usr/lib/tmac.x”.
- s* Stop prior to each page to permit paper loading. Printing is restarted by typing a ‘new-line’ character.
- h* Spaces are replaced where possible with tabs to speed up output (or reduce the size of the output file).
- q* Prompt names for insertions are not printed and the bell character is sent instead; the insertion is not echoed.
- i* Causes the standard input to be read after the files.

FILES

/usr/lib/suftab	suffix hyphenation tables
/tmp/rtn?	temporary
/usr/lib/tmac.?	standard macro files

SEE ALSO

NROFF User’s Manual (internal memorandum).
neqn (I)

BUGS

REQUEST REFERENCE AND INDEX

Request Form	Initial Value	If no Argument	Cause Break	Explanation
I. Page Control				
.pl +N	N=66	N=66	no	Page length.
.bp +N	N=1	-	yes	Begin page.
.pn +N	N=1	ignored	no	Page number.
.po +N	N=0	N=prev	no	Page offset.
.ne N	-	N=1	no	Need N lines.
.mk	none	-	no	Mark current line.
.rt	-	-	no	Return to marked line.
II. Text Filling, Adjusting, and Centering				
.br	-	-	yes	Break.
.fi	fill	-	yes	Fill output lines.
.nf	fill	-	yes	No filling and adjusting.
.ad c	adj,norm adjust	-	no	Adjust mode on.
.na	adjust	-	no	No adjusting.
.ce N	off	N=1	yes	Center N input text lines.
III. Line Spacing and Blank Lines				
.ls +N	N=1	N=prev	no	Line Spacing.
.sp N	-	N=1	yes	Space N lines
.lv N	-	N=1	no	Save N lines
.sv N	-	N=1	no	"
.os	-	-	no	Output Saved lines.
.ns	space	-	no	No-Space mode on.
.rs	-	-	no	Restore spacing.
.xh	off	-	no	Extra-half-line mode on.
IV. Line Length and Indenting				
.ll +N	N=65	N=prev	no	Line length.
.in +N	N=0	N=prev	yes	Indent.
.ti +N	-	N=1	yes	Temporary indent.
V. Macros, Diversion, and Line Traps				
.de xx	-	ignored	no	Define or redefine a macro.
.am xx	-	ignored	no	Append to a macro.
.ds xx	-	ignored	no	Define or redefine string.
.as xx	-	ignored	no	Append to a string.
.rm xx	-	-	no	Remove macro name.
.di xx	-	end	no	Divert output to macro "xx".
.da xx	-	end	no	Divert and append to "xx".
.wh -N xx	-	-	no	When; set a line trap.
.ch -N -M	-	-	no	Change trap line.
.ch xx -M	-	-	no	"
VI. Number Registers				
.nr ab +N -M -		no		Number Register.
.nr a +N -M -		no		"
.nc c	\n	\n	no	Number Character.
.ar	arabic	-	no	Arabic numbers.
.ro	arabic	-	no	Lower case roman numbers.
.RO	arabic	-	no	Upper case roman numbers.

VII. Input and Output Conventions and Character Translations

.ta N,M,...	-	none	no	Pseudotabs setting.
.tc c	space	space	no	Tab replacement character.
.lc c	.	.	no	Leader replacement character.
.ul n	-	N=1	no	Underline input text lines.
.cc c	.	.	no	Basic control character.
.c2 c	,	,	no	Nobreak control character.
.ec c	-	\	no	Escape character.
.li N	-	N=1	no	Accept input lines literally.
.tr abcd....	-	-	no	Translate on output.

VIII. Hyphenation.

.nh	on	-	no	No hyphenation.
.hy	on	-	no	Hyphenate.
.hc c	none	none	no	Hyphenation indicator character.

IX. Three Part Titles.

.tl 'left'center'right'	-	no	Title.	
.lt N	N=65	N=prev	no	Length of title.

X. Output Line Numbering.

.nm +N M S I	off	no	Number mode on or off, set parameters.
.np M S I -	reset	no	Number parameters set or reset.

XI. Conditional Input Line Acceptance

.if !N anything	-	no	If true accept line of "anything".
.if c anything-	no		"
.if !c anything	-	no	"
.if N anything	-	no	"

XII. Environment Switching.

.ev N	N=0	N=prev	no	Environment switched (pushed down).
-------	-----	--------	----	-------------------------------------

XIII. Insertions from the Standard Input Stream

.rd prompt	-	bell	no	Read insert.
.ex	-	-	no	Exit.

XIV. Input File Switching

.so filename	-	-	no	Switch source file (push down).
.nx filename	-	no		Next file.

XV. Miscellaneous

.tm mesg	-	-	no	Typewriter message
.ig	-	-	no	Ignore.
.fl	-	-	no	Flush output buffer.
.ab	-	-	no	Abort.

NAME

od - octal dump

SYNOPSIS

od [**-abcdho**] [*file*] [[**+**] offset[**.**][**b**]]

DESCRIPTION

Od dumps *file* in one or more formats as selected by the first argument. If the first argument is missing **-o** is default. The meanings of the format argument characters are:

- a** interprets words as PDP-11 instructions and dis-assembles the operation code. Unknown operation codes print as ???.
- b** interprets bytes in octal.
- c** interprets bytes in ascii. Unknown ascii characters are printed as \?.
- d** interprets words in decimal.
- h** interprets words in hex.
- o** interprets words in octal.

The *file* argument specifies which file is to be dumped. If no file argument is specified, the standard input is used. Thus *od* can be used as a filter.

The offset argument specifies the offset in the file where dumping is to commence. This argument is normally interpreted as octal bytes. If **.'** is appended, the offset is interpreted in decimal. If **'b'** is appended, the offset is interpreted in blocks. (A block is 512 bytes.) If the file argument is omitted, the offset argument must be preceded by **+**.

Dumping continues until end-of-file.

SEE ALSO

db (I)

BUGS

NAME

opr – off line print

SYNOPSIS

opr [—[*id*]] [-] [+] [+-]file ...

DESCRIPTION

Opr will arrange to have the 201 data phone daemon submit a job to the Honeywell 6070 to print the file arguments. Normally, the output appears at the GCOS central site. If the first argument is —, followed by an optional two-character remote station ID, the output is remoted to that station. If no station ID is given, R1 is assumed. (Station R1 has an IBM 1403 printer.)

Normally, each file is printed in the state it is found when the data phone daemon reads it. If a particular file argument is preceded by +, or a preceding argument of + has been encountered, then *opr* will make a copy for the daemon to print. If the file argument is preceded by -, or a preceding argument of - has been encountered, then *opr* will unlink (remove) the file.

If there are no arguments except for the optional —, then the standard input is read and off-line printed. Thus *opr* may be used as a filter.

FILES

/usr/dpd/*	spool area
/etc/passwd	personal ident cards
/etc/dpd	daemon

SEE ALSO

dpd (VIII), passwd (V)

BUGS

NAME

passwd – set login password

SYNOPSIS

passwd name password

DESCRIPTION

The *password* is placed on the given login name. This can only be done by the person corresponding to the login name or by the super-user. An explicit null argument ("") for the password argument will remove any password from the login name.

FILES

/etc/passwd

SEE ALSO

login(I), passwd(V), crypt(III)

BUGS

NAME

pfe – print floating exception

SYNOPSIS

pfe

DESCRIPTION

Pfe examines the floating point exception register and print a diagnostic for the last floating point exception.

SEE ALSO

signal (II)

BUGS

Since the system does not save the exception register in a core image file, the message refers to the last error encountered by anyone. Floating exceptions are therefore volatile.

NAME

pr - print file

SYNOPSIS

pr [**-h** *header*] [**-n**] [**+n**] [**-wn**] [**-ln**] [**-t**] [*name* . . .]

DESCRIPTION

Pr produces a printed listing of one or more files. The output is separated into pages headed by a date, the name of the file or a header (if any), and the page number. If there are no file arguments, *pr* prints its standard input, and is thus usable as a filter.

Options apply to all following files but may be reset between files:

- n** produce *n*-column output
- +n** begin printing with page *n*
- h** treat the next argument as a header
- wn** for purposes of multi-column output, take the width of the page to be *n* characters instead of the default 72
- ln** take the length of the page to be *n* lines instead of the default 66
- t** do not print the 5-line header or the 5-line trailer normally supplied for each page

If there is a header in force, it is printed in place of the file name. Interconsole messages via write(I) are forbidden during a *pr*.

FILES

/dev/tty? to suspend messages.

SEE ALSO

cat(I), cp(I)

DIAGNOSTICS

none; files not found are ignored

BUGS

NAME

prof – display profile data

SYNOPSIS

prof [**-v**] [**-a**] [**-l**] [file]

DESCRIPTION

Prof interprets the file *mon.out* produced by the *monitor* subroutine. Under default modes, the symbol table in the named object file (*a.out* default) is read and correlated with the *mon.out* profile file. For each external symbol, the percentage of time spent executing between that symbol and the next is printed (in decreasing order), together with the number of times that routine was called and the number of milliseconds per call.

If the **-a** option is used, all symbols are reported rather than just external symbols. If the **-l** option is used, the output is listed by symbol value rather than decreasing percentage. If the **-v** option is used, all printing is suppressed and a profile plot is produced on the 611 display.

In order for the number of calls to a routine to be tallied, the **-p** option of *cc* must have been given when the file containing the routine was compiled.

FILES

mon.out for profile
a.out for namelist
/dev/vt0 for plotting

SEE ALSO

monitor(III), profil(II), cc(I)

BUGS

NAME

ps – process status

SYNOPSIS

ps [**aklx**] [namelist]

DESCRIPTION

Ps prints certain indicia about active processes. The **a** flag asks for information about all processes with teletypes (ordinarily only one's own processes are displayed); **x** asks even about processes with no typewriter; **l** asks for a long listing. Ordinarily only the typewriter number (if not one's own), the process number, and an approximation to the command line are given. If the **k** flag is specified, the special file */usr/sys/core* is used in place of */dev/mem*. This is used for post-mortem system debugging.

The long listing is columnar and contains

A number encoding the state (last digit) and flags (first 1 or 2 digits) of the process.

The priority of the process; high numbers mean low priority.

A number related in some unknown way to the scheduling heuristic.

The last character of the control typewriter of the process.

The process unique number (as in certain cults it is possible to kill a process if you know its true name).

The size in blocks of the core image of the process.

The last column if non-blank tells the core address in the system of the event which the process is waiting for; if blank, the process is running.

Ps makes an educated guess as to the file name and arguments given when the process was created by examining core memory or the swap area. The method is inherently somewhat unreliable and in any event a process is entitled to destroy this information, so the names cannot be counted on too much.

FILES

/unix system namelist
/dev/mem core memory
/dev/rf0 swap device
/dev/rk0 optional mem file

SEE ALSO

kill (I)

BUGS

The command has assumptions built into it about the number of typewriters that exist and what hardware is used to interface them. It also has built into it the name of the device used for swapping.

NAME

pwd – working directory name

SYNOPSIS

pwd

DESCRIPTION

Pwd prints the pathname of the working (current) directory.

SEE ALSO

chdir (I)

BUGS

The algorithm sometimes fails on crossing mounted file system.

NAME

rew – rewind tape

SYNOPSIS

rew [[**m**]digit]

DESCRIPTION

Rew rewinds DECtape or magtape drives. The digit is the logical tape number, and should range from 0 to 7. if the digit is preceded by **m**, *rew* applies to magtape rather than DECtape. A missing digit indicates drive 0.

FILES

/dev/tap?
/dev/mt?

BUGS

NAME

rm – remove (unlink) files

SYNOPSIS

rm [**-f**] [**-r**] name ...

DESCRIPTION

Rm removes the entries for one or more files from a directory. If an entry was the last link to the file, the file is destroyed. Removal of a file requires write permission in its directory, but neither read nor write permission on the file itself.

If there is no write permission to a file designated to be removed, *rm* will print the file name, its mode and then read a line from the standard input. If the line begins with **y**, the file is removed, otherwise it is not. The optional argument **-f** prevents this interaction.

If a designated file is a directory, an error comment is printed unless the optional argument **-r** has been used. In that case, *rm* recursively deletes the entire contents of the specified directory. To remove directories *per se* see *rmdir*(I).

FILES

/etc/glob to implement the **-r** flag

SEE ALSO

rmdir(I)

BUGS

When *rm* removes the contents of a directory under the **-r** flag, full pathnames are not printed in diagnostics.

NAME

`rmdir` – remove directory

SYNOPSIS

`rmdir` `dir ...`

DESCRIPTION

Rmdir removes (deletes) directories. The directory must be empty (except for the standard entries `‘.’` and `‘..’`, which *rmdir* itself removes). Write permission is required in the directory in which the directory appears.

BUGS

Needs a `–r` flag. Actually, write permission in the directory’s parent is *not* required.

NAME

roff - format text

SYNOPSIS

roff [*+n*] [*-n*] [*-s*] [*-h*] file ...

DESCRIPTION

Roff formats text according to control lines embedded in the text in the given files. Encountering a nonexistent file terminates printing. Incoming interconsole messages are turned off during printing. The optional flag arguments mean:

- +n* Start printing at the first page with number *n*.
- n* Stop printing at the first page numbered higher than *n*.
- s* Stop before each page (including the first) to allow paper manipulation; resume on receipt of an interrupt signal.
- h* Insert tabs in the output stream to replace spaces whenever appropriate.

A Request Summary is attached.

FILES

/usr/lib/suftabsuffix hyphenation tables
/tmp/rtn?temporary

SEE ALSO

nroff (I), troff (I)

BUGS

Roff is the simplest of the runoff programs, but is virtually undocumented.

REQUEST SUMMARY

<i>Request</i>	<i>Break</i>	<i>Initial</i>	<i>Meaning</i>
.ad	yes	yes	Begin adjusting right margins.
.ar	no	arabic	Arabic page numbers.
.br	yes	-	Causes a line break – the filling of the current line is stopped.
.bl n	yes	-	Insert of n blank lines, on new page if necessary.
.bp +n	yes	n=1	Begin new page and number it n; no n means '+1'.
.cc c	no	c=.	Control character becomes 'c'.
.ce n	yes	-	Center the next n input lines, without filling.
.de xx	no	-	Define macro named 'xx' (definition ends on line beginning '..').
.ds	yes	no	Double space; same as '.ls 2'.
.ef t	no	t='''	Even foot title becomes t.
.eh t	no	t='''	Even head title becomes t.
.fi	yes	yes	Begin filling output lines.
.fo	no	t='''	All foot titles are t.
.hc c	no	none	Hyphenation character set to 'c'.
.he t	no	t='''	All head titles are t.
.hx	no	-	Title lines are suppressed.
.hy n	no	n=1	Hyphenation is done, if n=1; and is not done, if n=0.
.ig	no	-	Ignore input lines through a line beginning with '..'.
.in +n	yes	-	Indent n spaces from left margin.
.ix +n	no	-	Same as '.in' but without break.
.li n	no	-	Literal, treat next n lines as text.
.ll +n	no	n=65	Line length including indent is n characters.
.ls +n	yes	n=1	Line spacing set to n lines per output line.
.m1 n	no	n=2	Put n blank lines between the top of page and head title.
.m2 n	no	n=2	n blank lines put between head title and beginning of text on page.
.m3 n	no	n=1	n blank lines put between end of text and foot title.
.m4 n	no	n=3	n blank lines put between the foot title and the bottom of page.
.na	yes	no	Stop adjusting the right margin.
.ne n	no	-	Begin new page, if n output lines cannot fit on present page.
.nn +n	no	-	The next n output lines are not numbered.
.n1	no	no	Number output lines; start with 1 each page
.n2 n	no	no	Number output lines; stop numbering if n=0.
.ni +n	no	n=0	Line numbers are indented n.
.nf	yes	no	Stop filling output lines.
.nx filename	-	Change to input file 'filename'.	
.of t	no	t='''	Odd foot title becomes t.
.oh t	no	t='''	Odd head title becomes t.
.pa +n	yes	n=1	Same as '.bp'.
.pl +n	no	n=66	Total paper length taken to be n lines.
.po +n	no	n=0	Page offset. All lines are preceded by N spaces.
.ro	no	arabic	Roman page numbers.
.sk n	no	-	Produce n blank pages starting next page.
.sp n	yes	-	Insert block of n blank lines.
.ss	yes	yes	Single space output lines, equivalent to '.ls 1'.
.ta N M ...	-	-	Pseudotab settings. Initial tab settings are columns 9,17,25,...
.tc c	no	c=' '	Tab replacement character becomes 'c'.
.ti +n	yes	-	Temporarily indent next output line n space.
.tr abcd..	no	-	Translate a into b, c into d, etc.
.ul n	no	-	Underline the letters and numbers in the next n input lines.

NAME

sh – shell (command interpreter)

SYNOPSIS

sh [name [arg1 ... [arg9]]]

DESCRIPTION

Sh is the standard command interpreter. It is the program which reads and arranges the execution of the command lines typed by most users. It may itself be called as a command to interpret files of commands. Before discussing the arguments to the Shell used as a command, the structure of command lines themselves will be given.

Commands. Each command is a sequence of non-blank command arguments separated by blanks. The first argument specifies the name of a command to be executed. Except for certain types of special arguments discussed below, the arguments other than the command name are passed without interpretation to the invoked command.

If the first argument is the name of an executable file, it is invoked; otherwise the string *‘/bin/’* is prepended to the argument. (In this way most standard commands, which reside in *‘/bin/’*, are found.) If no such command is found, the string *‘/usr/’* is further prepended (to give *‘/usr/bin/command’*) and another attempt is made to execute the resulting file. (Certain lesser-used commands live in *‘/usr/bin/’*.) If the *‘/usr/bin/’* file exists, but is not executable, it is used by the Shell as a command file. That is to say it is executed as though it were typed from the console. If all attempts fail, a diagnostic is printed.

Command lines. One or more commands separated by *‘|’* or *‘^’* constitute a *pipeline*. The standard output of each command but the last in a pipeline is taken as the standard input of the next command. Each command is run as a separate process, connected by pipes (see *pipe(II)*) to its neighbors. A command line contained in parentheses *‘()’* may appear in place of a simple command as an element of a pipeline.

A *command line* consists of one or more pipelines separated, and perhaps terminated by *‘;’* or *‘&’*. The semicolon designates sequential execution. The ampersand causes the preceding pipeline to be executed without waiting for it to finish. The process id of such a pipeline is reported, so that it may be used if necessary for a subsequent *wait* or *kill*.

Termination Reporting. If a command (not followed by *‘&’*) terminates abnormally, a message is printed. (All terminations other than exit and interrupt are considered abnormal.) Termination reports for commands followed by *‘&’* are given upon receipt of the first command subsequent to the termination of the command, or when a *wait* is executed. The following is a list of the abnormal termination messages:

- Bus error
- Trace/BPT trap
- Illegal instruction
- IOT trap
- EMT trap
- Bad system call
- Quit
- Floating exception
- Memory violation
- Killed

If a core image is produced, *‘– Core dumped’* is appended to the appropriate message.

Redirection of I/O. There are three character sequences that cause the immediately following string to be interpreted as a special argument to the Shell itself. Such an argument may appear anywhere among the arguments of a simple command, or before or after a parenthesized command list, and is associated with that command or command list.

An argument of the form '<arg' causes the file 'arg' to be used as the standard input (file descriptor 0) of the associated command.

An argument of the form '>arg' causes file 'arg' to be used as the standard output (file descriptor 1) for the associated command. 'Arg' is created if it did not exist, and in any case is truncated at the outset.

An argument of the form '>>arg' causes file 'arg' to be used as the standard output for the associated command. If 'arg' did not exist, it is created; if it did exist, the command output is appended to the file.

For example, either of the command lines

```
ls >junk; cat tail >>junk
( ls; cat tail ) >junk
```

creates, on file 'junk', a listing of the working directory, followed immediately by the contents of file 'tail'.

Either of the constructs '>arg' or '>>arg' associated with any but the last command of a pipeline is ineffectual, as is '<arg' in any but the first.

In commands called by the Shell, file descriptor 2 refers to the standard output of the Shell before any redirection. Thus filters may write diagnostics to a location where they have a chance to be seen.

Generation of argument lists. If any argument contains any of the characters '?', '*' or '[', it is treated specially as follows. The current directory is searched for files which *match* the given argument.

The character '*' in an argument matches any string of characters in a file name (including the null string).

The character '?' matches any single character in a file name.

Square brackets '['...']' specify a class of characters which matches any single file-name character in the class. Within the brackets, each ordinary character is taken to be a member of the class. A pair of characters separated by '-' places in the class each character lexically greater than or equal to the first and less than or equal to the second member of the pair.

Other characters match only the same character in the file name.

For example, '*' matches all file names; '?' matches all one-character file names; '[ab]*.s' matches all file names beginning with 'a' or 'b' and ending with '.s'; '?[zi-m]' matches all two-character file names ending with 'z' or the letters 'i' through 'm'.

If the argument with '*' or '?' also contains a '/', a slightly different procedure is used: instead of the current directory, the directory used is the one obtained by taking the argument up to the last '/' before a '*' or '?'. The matching process matches the remainder of the argument after this '/' against the files in the derived directory. For example: '/usr/dmr/a*.s' matches all files in directory '/usr/dmr' which begin with 'a' and end with '.s'.

In any event, a list of names is obtained which match the argument. This list is sorted into alphabetical order, and the resulting sequence of arguments replaces the single argument containing the '*', '[', or '?'. The same process is carried out for each argument (the resulting lists are *not* merged) and finally the command is called with the resulting list of arguments.

For example: directory /usr/dmr contains the files a1.s, a2.s, ..., a9.s. From any directory, the command

```
as /usr/dmr/a?.s
```

calls *as* with arguments /usr/dmr/a1.s, /usr/dmr/a2.s, ... /usr/dmr/a9.s in that order.

Quoting. The character '\' causes the immediately following character to lose any special meaning it may have to the Shell; in this way '<', '>', and other characters meaningful to the Shell may be passed as part of arguments. A special case of this feature allows the continuation of commands onto more than one line: a new-line preceded by '\' is translated into a blank.

Sequences of characters enclosed in double (") or single (') quotes are also taken literally. For example:

```
ls | pr -h "My directory"
```

causes a directory listing to be produced by *ls*, and passed on to *pr* to be printed with the heading 'My directory'. Quotes permit the inclusion of blanks in the heading, which is a single argument to *pr*.

Argument passing. When the Shell is invoked as a command, it has additional string processing capabilities. Recall that the form in which the Shell is invoked is

```
sh [ name [ arg1 ... [ arg9 ] ] ]
```

The *name* is the name of a file which will be read and interpreted. If not given, this subinstance of the Shell will continue to read the standard input file.

In command lines in the file (not in command input), character sequences of the form '\$n', where *n* is a digit, are replaced by the *n*th argument to the invocation of the Shell (argn). '\$0' is replaced by *name*.

End of file. An end-of-file in the Shell's input causes it to exit. A side effect of this fact means that the way to log out from UNIX is to type an EOT.

Special commands. The following commands are treated specially by the Shell.

chdir is done without spawning a new process by executing *sys chdir* (II).

login is done by executing */bin/login* without creating a new process.

wait is done without spawning a new process by executing *sys wait* (II).

shift is done by manipulating the arguments to the Shell.

'.' is simply ignored.

Command file errors; interrupts. Any Shell-detected error, or an interrupt signal, during the execution of a command file causes the Shell to cease execution of that file.

Processes that are created with a '&' ignore interrupts. Also if such a process has not redirected its input with a '<', its input is automatically redirected to the zero length file */dev/null*.

FILES

/etc/glob, which interprets '*', '?', and '['.
/dev/null as a source of end-of-file.

SEE ALSO

'The UNIX Time-sharing System', which gives the theory of operation of the Shell.
chdir (I), *login* (I), *wait* (I), *shift* (I)

BUGS

There is no way to redirect the diagnostic output.

NAME

shift – adjust Shell arguments

SYNOPSIS

shift

DESCRIPTION

Shift is used in Shell command files to shift the argument list left by 1, so that old **\$2** can now be referred to by **\$1** and so forth. *Shift* is useful to iterate over several arguments to a command file. For example, the command file

```
: loop
if $1x = x exit
pr -3 $1
shift
goto loop
```

prints each of its arguments in 3-column format.

Shift is executed within the Shell.

SEE ALSO

sh (1)

BUGS

SIZE (I)

9/2/72

SIZE (I)

NAME

size – size of an object file

SYNOPSIS

size [object ...]

DESCRIPTION

Size prints the (decimal) number of bytes required by the text, data, and bss portions, and their sum in octal and decimal, of each object-file argument. If no file is specified, **a.out** is used.

BUGS

NAME

sleep – suspend execution for an interval

SYNOPSIS

sleep time

DESCRIPTION

Sleep will suspend execution for *time* seconds. It is used to execute a command in a certain amount of time as in:

(sleep 105; command)&

Or to execute a command every so often as in this shell command file:

```
: loop
command
sleep 37
goto loop
```

SEE ALSO

sleep(II)

BUGS

Time must be less than 65536 seconds.

NAME

sort – sort or merge files

SYNOPSIS

sort [**-abdnrtx**] [*+pos* [*-pos*]] . . . [**-mo**] [name] . . .

DESCRIPTION

Sort sorts all the named files together and writes the result on the standard output. The name ‘-’ means the standard input. The standard input is also used if no input file names are given. Thus *sort* may be used as a filter.

The default sort key is an entire line. Default ordering is lexicographic in ASCII collating sequence, except that lower-case letters are considered the same as the corresponding upper-case letters. Non-ASCII bytes are ignored. The ordering is affected by the flags **-abdnrt**, one or more of which may appear:

- a** Do not map lower case letters.
- b** Leading blanks (spaces and tabs) are not included in fields.
- d** ‘Dictionary’ order: only letters, digits and blanks are significant in ASCII comparisons.
- n** An initial numeric string, consisting of optional minus sign, digits and optionally included decimal point, is sorted by arithmetic value.
- r** Reverse the sense of comparisons.
- tx** Tab character between fields is *x*.

Selected parts of the line, specified by *+pos* and *-pos*, may be used as sort keys. *Pos* has the form *m.n*, where *m* specifies a number of fields to skip, and *n* a number of characters to skip further into the next field. A missing is taken to be 0. *+pos* denotes the beginning of the key; *-pos* denotes the first position after the key (end of line by default). The ordering rule may be overridden for a particular key by appending one or more of the flags **abdnr** to *+pos*.

When no tab character has been specified, a field consists of nonblanks and any preceding blanks. Under the **-b** flag, leading blanks are excluded from a field. When a tab character has been specified, a field is a string ending with a tab character.

When keys are specified, later keys are compared only when all earlier ones compare equal. Lines that compare equal are ordered with all bytes significant.

These flag arguments are also understood:

- m** Merge only, the input files are already sorted.
- o** The next argument is the name of an output file to use instead of the standard output. This file may be the same as one of the inputs, except under the merge flag **-m**.

FILES

/usr/tmp/stm???

NAME

spell – find spelling errors

SYNOPSIS

spell file ...

DESCRIPTION

Spell attacks the same problem as *typo* (I), but from the opposite direction. It extracts words from the input files and looks them up in *Webster's Seventh Collegiate Dictionary*; any words which appear neither in the dictionary nor in a list of about 2000 words frequently occurring in Bell Laboratories documents are listed on the output file *sp.out*. Words which are reasonable transformations of dictionary entries (e.g. a dictionary entry plus *s*) are so marked; words which could not be found even when transformed are marked with asterisks.

The process takes on the order of 5 to 10 minutes. There is a limit of nine input files.

FILES

/usr/lib/w2006, /usr/dict/words, sp.out; spjnkq[123] are temporaries.

SEE ALSO

typo (I)

BUGS

There should be no limit on the number of input files.

More suffixes, and perhaps some prefixes, should be added.

It should be usable as a filter.

NAME

split – split a file into pieces

SYNOPSIS

split *-n* [file [name]]

DESCRIPTION

Split reads *file* and writes it in *n*-line pieces (default 1000), as many as necessary, onto a set of output files. The name of the first output file is *name* with **aa** appended, and so on lexicographically. If no output name is given, **x** is default.

If no input file is given, or if *-* is given in its stead, then the standard input file is used.

BUGS

NAME

strip — remove symbols and relocation bits

SYNOPSIS

strip name ...

DESCRIPTION

Strip removes the symbol table and relocation bits ordinarily attached to the output of the assembler and loader. This is useful to save space after a program has been debugged.

The effect of *strip* is the the same as use of the **-s** option of *ld*.

FILES

/tmp/stm? temporary file

SEE ALSO

ld(I), as(I)

BUGS

NAME

stty – set teletype options

SYNOPSIS

stty option ...

DESCRIPTION

Stty will set certain I/O options on the current output teletype. The option strings are selected from the following set:

even	allow even parity
–even	disallow even parity
odd	allow odd parity
–odd	disallow odd parity
raw	raw mode input (no erase, kill, interrupt, quit, EOT; parity bit passed back)
–raw	negate raw mode
–nl	allow carriage return for new-line, and output CR-LF for carriage return or new-line
nl	accept only new-line to end lines
echo	echo back every character typed
–echo	do not echo characters
lcase	map upper case to lower case
–lcase	do not map case
–tabs	replace tabs by spaces in output
tabs	preserve tabs
delay	calculate cr, tab, and form-feed delays
–delay	no cr/tab/ff delays
tdelay	calculate tab delays
–tdelay	no tab delays

SEE ALSO

stty (II)

BUGS

There should be ‘package’ options such as **execuport**, **33**, or **terminet**.

SUM(I)

3/15/72

SUM(I)

NAME

sum - sum file

SYNOPSIS

sum name ...

DESCRIPTION

Sum sums the contents of the bytes (mod 2^{16}) of one or more files and prints the answer in octal. A separate sum is printed for each file specified, along with the number of whole or partial 512-byte blocks read.

In practice, *sum* is often used to verify that all of a special file can be read without error.

BUGS

TEE (I)

3/6/74

TEE (I)

NAME

tee – pipe fitting

SYNOPSIS

tee [name ...]

DESCRIPTION

Tee transcribes the standard input to the standard output and makes copies in the named files.

BUGS

NAME

time – time a command

SYNOPSIS

time command

DESCRIPTION

The given command is executed; after it is complete, *time* prints the elapsed time during the command, the time spent in the system, and the time spent in execution of the command.

The execution time can depend on what kind of memory the program happens to land in; the user time in MOS is often half what it is in core.

The times are printed on the diagnostic output stream.

BUGS

Elapsed time is accurate to the second, while the CPU times are measured to the 60th second. Thus the sum of the CPU times can be up to a second larger than the elapsed time.

NAME

tp – manipulate DECTape and magtape

SYNOPSIS

tp [key] [name ...]

DESCRIPTION

tp saves and restores selected portions of the file system hierarchy on DECTape or mag tape. Its actions are controlled by the *key* argument. The key is a string of characters containing at most one function letter and possibly one or more function modifiers. Other arguments to the command are file or directory names specifying which files are to be dumped, restored, or listed.

The function portion of the key is specified by one of the following letters:

- r** The indicated files and directories, together with all subdirectories, are dumped onto the tape. If files with the same names already exist, they are replaced. 'Same' is determined by string comparison, so './abc' can never be the same as '/usr/dmr/abc' even if '/usr/dmr' is the current directory. If no file argument is given, '.' is the default.
- u** updates the tape. **u** is the same as **r**, but a file is replaced only if its modification date is later than the date stored on the tape; that is to say, if it has changed since it was dumped. **u** is the default command if none is given.
- d** deletes the named files and directories from the tape. At least one file argument must be given. This function is not permitted on magtapes.
- x** extracts the named files from the tape to the file system. The owner, mode, and date-modified are restored to what they were when the file was dumped. If no file argument is given, the entire contents of the tape are extracted.
- t** lists the names of all files stored on the tape which are the same as or are hierarchically below the file arguments. If no file argument is given, the entire contents of the tape is listed.

The following characters may be used in addition to the letter which selects the function desired.

- m** Specifies magtape as opposed to DECTape.
- 0,...,7** This modifier selects the drive on which the tape is mounted. For DECTape, 'x' is default; for magtape '0' is the default.
- v** Normally *tp* does its work silently. The **v** (verbose) option causes it to type the name of each file it treats preceded by the function letter. With the **t** function, **v** gives more information about the tape entries than just the name.
- c** means a fresh dump is being created; the tape directory will be zeroed before beginning. Usable only with **r** and **u**. This option is assumed with magtape since it is impossible to selectively overwrite magtape.
- f** causes new entries on tape to be 'fake' in that no data is present for these entries. Such fake entries cannot be extracted. Usable only with **r** and **u**.
- i** Errors reading and writing the tape are noted, but no action is taken. Normally, errors cause a return to the command level.
- w** causes *tp* to pause before treating each file, type the indicative letter and the file name (as with **v**) and await the user's response. Response **y** means 'yes', so the file is treated. Null response means 'no', and the file does not take part in whatever is being done. Response **x** means 'exit'; the *tp* command terminates immediately. In the **x** function, files previously asked about have been extracted already. With **r**, **u**, and **d** no change has been made to the tape.

FILES

/dev/tap?
/dev/mt?

TP (I)

10/15/73

TP (I)

DIAGNOSTICS

Several; the non-obvious one is 'Phase error', which means the file changed after it was selected for dumping but before it was dumped.

BUGS

NAME

tr – transliterate

SYNOPSIS

tr [**-cds**] [string1 [string2]]

DESCRIPTION

Tr copies the standard input to the standard output with substitution or deletion of selected characters. Input characters found in *string1* are mapped into the corresponding characters of *string2*. If *string2* is short, it is padded with corresponding characters from *string1*. Any combination of the options **-cds** may be used. **-c** complements the set of characters in *string1* with respect to the universe of characters whose ascii codes are 001 through 377 octal. **-d** deletes all input characters in *string1*. **-s** squeezes all strings of repeated output characters that are in *string2* to single characters.

The following abbreviation conventions may be used to introduce ranges of characters or repeated characters into the strings:

[*a-b*] stands for the string of characters whose ascii codes run from character *a* to character *b*.

[*a*n*], where *n* is an integer or empty, stands for *n*-fold repetition of character *a*. *n* is taken to be octal or decimal according as its first digit is or is not zero. A zero or missing *n* is taken to be huge; this facility is useful for padding *string2*.

The escape character '\ ' may be used as in *sh* to remove special meaning from any character in a string. In addition, '\ ' followed by 1, 2 or 3 octal digits stands for the character whose ascii code is given by those digits.

The following example creates a list of all the words in 'file1' one per line in 'file2', where a word is taken to be a maximal string of alphabetic. The strings are quoted to protect the special characters from interpretation by the Shell; 012 is the ascii code for newline.

```
tr -cs "[A-Z][a-z]" "[\012*]" <file1 >file2
```

SEE ALSO

sh (I), ed (I), ascii (VII)

BUGS

Won't handle ascii NUL in *string1* or *string2*; always deletes NUL from input.

NAME

troff – format text

SYNOPSIS

troff [*+n*] [*-n*] [*-nn*] [*-mx*] [*-t*] [*-f*] [*-w*] [*-i*] [*-a*] [*-pn*] files

DESCRIPTION

Troff formats text for a Graphic Systems phototypesetter according to control lines embedded in the text files. It reads the standard input if no file arguments are given. The non-file option arguments are interpreted as follows:

- +n* Commence typesetting at the first page numbered *n* or larger.
- n* Stop after page *n*.
- nn* First generated (not necessarily printed) page is given the number *n*; simulates “.pn *n*”.
- mx* Prepends a standard macro file; simulates “.so /usr/lib/tmac.x”.
- t* Place output on standard output instead of the phototypesetter.
- f* Refrain from feeding out paper and stopping the phototypesetter at the end.
- w* Wait until phototypesetter is available, if currently busy.
- i* Read from standard input after the files have been exhausted.
- a* Send a printable approximation of the results to the standard output.
- pn* Print all characters with point-size *n* while retaining all prescribed spacings and motions.

FILES

/usr/lib/suftab	suffix hyphenation tables
/tmp/rtn?	temporary
/usr/lib/tmac.x	standard macro files

SEE ALSO

TROFF User's Manual (internal memorandum).
TROFF Made Trivial (internal memorandum).
nroff (I), eqn (I), catsim (VI)

BUGS

NAME

tss – interface to MH-TSS

SYNOPSIS

tss

DESCRIPTION

Tss will call the Honeywell 6070 on the 201 data phone. It will then go into direct access with MH-TSS. Output generated by MH-TSS is typed on the standard output and input requested by MH-TSS is read from the standard input with UNIX typing conventions.

An interrupt signal is transmitted as a ‘break’ to MH-TSS.

Input lines beginning with ‘!’ are interpreted as UNIX commands. Input lines beginning with ‘~’ are interpreted as commands to the interface routine.

~<file	insert input from named UNIX file
~>file	deliver tss output to named UNIX file
~p	pop the output file
~q	disconnect from tss (quit)
~r file	receive from HIS routine csr/daccopy
~s file	send file to HIS routine csr/daccopy

Ascii files may be most efficiently transmitted using the HIS routine csr/daccopy in this fashion. Bold face text comes from MH-TSS. *Aftname* is the 6070 file to be dealt with; *file* is the UNIX file.

SYSTEM? csr/daccopy (s) *aftname*

Send Encoded File ~s *file*

SYSTEM? csr/daccopy (r) *aftname*

Receive Encoded File ~r *file*

FILES

/dev/dn0, /dev/dp0, /etc/msh

DIAGNOSTICS

Most often, ‘Transmission error on last message.’

BUGS

When problems occur, and they often do, *tss* exits rather abruptly.

TTY (I)

3/15/72

TTY (I)

NAME

tty – get typewriter name

SYNOPSIS

tty

DESCRIPTION

Tty gives the name of the user's typewriter in the form 'ttn' for *n* a digit or letter. The actual path name is then '/dev/ttn'.

DIAGNOSTICS

'not a tty' if the standard input file is not a typewriter.

BUGS

NAME

typo – find possible typos

SYNOPSIS

typo [**-1**] [**-n**] file ...

DESCRIPTION

Typo hunts through a document for unusual words, typographic errors, and *hapax legomena* and prints them on the standard output.

The words used in the document are printed out in decreasing order of peculiarity along with an index of peculiarity. An index of 10 or more is considered peculiar. Printing of certain very common English words is suppressed.

The statistics for judging words are taken from the document itself, with some help from known statistics of English. The **-n** option suppresses the help from English and should be used if the document is written in, for example, Urdu.

The **-1** option causes the final output to appear in a single column instead of three columns. The normal header and pagination is also suppressed.

Roff (I) and nroff (I) control lines are ignored. Upper case is mapped into lower case. Quote marks, vertical bars, hyphens, and ampersands are stripped from within words. Words hyphenated across lines are put back together.

FILES

/tmp/ttmp??
/usr/lib/salt
/usr/lib/w2006

BUGS

Because of the mapping into lower case and the stripping of special characters, words may be hard to locate in the original text.

The escape sequences of troff (I) are not correctly recognized.

NAME

uniq – report repeated lines in a file

SYNOPSIS

uniq [**-udc** [**+n**] [**-n**]] [input [output]]

DESCRIPTION

Uniq reads the input file comparing adjacent lines. In the normal case, the second and succeeding copies of repeated lines are removed; the remainder is written on the output file. Note that repeated lines must be adjacent in order to be found; see *sort*(I). If the **-u** flag is used, just the lines that are not repeated in the original file are output. The **-d** option specifies that one copy of just the repeated lines is to be written. The normal mode output is the union of the **-u** and **-d** mode outputs.

The **-c** option supersedes **-u** and **-d** and generates an output report in default style but with each line preceded by a count of the number of times it occurred.

The *n* arguments specify skipping an initial portion of each line in the comparison:

- n** The first *n* fields together with any blanks before each are ignored. A field is defined as a string of non-space, non-tab characters separated by tabs and spaces from its neighbors.
- +n** The first *n* characters are ignored. Fields are skipped before characters.

SEE ALSO

sort (I), *comm* (I)

BUGS

NAME

wait – await completion of process

SYNOPSIS

wait

DESCRIPTION

Wait until all processes started with **&** have completed, and report on abnormal terminations.

Because *sys wait* must be executed in the parent process, the Shell itself executes *wait*, without creating a new process.

SEE ALSO

sh (I)

BUGS

After executing *wait* you are committed to waiting until termination, because interrupts and quits are ignored by all processes concerned. The only out, if the process does not terminate, is to *kill* it from another terminal or to hang up.

NAME

wc – word count

SYNOPSIS

wc [**-rlwapc**] [name ...]

DESCRIPTION

Wc counts lines and words in the named files, or in the standard input if no name appears. A word is a maximal string of ascii graphics delimited by spaces, tabs or newlines. Other characters are always ignored.

Any of the following options may appear in any order. When any option other than **-r** appears, only the specified information is reported in the order in which the options occur. Otherwise, printing is as with option **-lw**.

- r** Ignore all *roff*, *nroff* and *troff* control lines, i.e. lines beginnngin with ‘.’ or ‘’.
- l** Print count of lines
- w** Print count of words
- a** Print count of alphanumeric strings, with underscore taken as alphanumeric.
- p** Print count of punctuation strings, i.e. all strings of printing characters other than alphanumerics.
- c** Print count of *roff* control lines, regardless of **-r**.

DIAGNOSTICS

‘Cannot open’ for unopenable file.

BUGS

NAME

who – who is on the system

SYNOPSIS

who [who-file]

DESCRIPTION

Who, without an argument, lists the name, typewriter channel, and login time for each current UNIX user.

Without an argument, *who* examines the */tmp/utmp* file to obtain its information. If a file is given, that file is examined. Typically the given file will be */tmp/wtmp*, which contains a record of all the logins since it was created. Then *who* will list logins, logouts, and crashes since the creation of the *wtmp* file.

Each login is listed with user name, typewriter name (with ‘/dev/’ suppressed), and date and time. When an argument is given, logouts produce a similar line without a user name. Reboots produce a line with ‘x’ in the place of the device name, and a fossil time indicative of when the system went down.

FILES

/tmp/utmp

SEE ALSO

login (I), init (VII)

BUGS

NAME

write – write to another user

SYNOPSIS

write user [ttyno]

DESCRIPTION

Write copies lines from your typewriter to that of another user. When first called, it sends the message

message from yourname...

The recipient of the message should write back at this point. Communication continues until an end of file is read from the typewriter or an interrupt is sent. At that point *write* writes 'EOT' on the other terminal and exits.

If you want to write to a user who is logged in more than once, the *ttyno* argument may be used to indicate the last character of the appropriate typewriter name.

Permission to write may be denied or granted by use of the *mesg* command. At the outset writing is allowed. Certain commands, in particular *roff* and *pr*, disallow messages in order to prevent messy output.

If the character '!' is found at the beginning of a line, *write* calls the mini-shell *msh* to execute the rest of the line as a command.

The following protocol is suggested for using *write*: when you first write to another user, wait for him to write back before starting to send. Each party should end each message with a distinctive signal ((o) for 'over' is conventional) that the other may reply. (oo) (for 'over and out') is suggested when conversation is about to be terminated.

FILES

/tmp/utmp to find user
/etc/msh to execute '!'

SEE ALSO

mesg (I), who (I)

BUGS