



Documents for the PWB/UNIX Time-Sharing System

Edition 1.0

T. A. Dolotta
R. C. Haight
E. M. Piskorik

Editors

October 1977

The enclosed PWB/UNIX documentation is supplied in accordance with the Software Agreement you have with the Western Electric Company.

Bell Telephone Laboratories, Incorporated

UNIX is a Trademark of Bell Laboratories.

The enclosed documents were set on a Graphic Systems, Inc. phototypesetter driven by the TROFF formatting program. Their text was prepared using the ED text editor.

Documents for the PWB/UNIX Time-Sharing System

Annotated Table of Contents

Each item carries the date of its latest revision; most items also give the number of their last page.

G. General:

- G.1 *PWB/UNIX—Overview and Synopsis of Facilities* (6/77)
T. A. Dolotta and R. C. Haight (p. 18)
Summarizes the salient features of Programmer's Workbench/UNIX, a program development and text processing facility.
- G.2 *The UNIX Time-Sharing System* (7/74)
D. M. Ritchie and K. Thompson (p. 16)
Reprinted from *Comm. ACM*. Good overview of UNIX, but written long ago.
- G.3 *The UNIX Time-sharing System—A Retrospective* (1/77)
D. M. Ritchie (p. 14)
A more recent discussion.
- G.4 *PWB/UNIX Papers from the Second Intern. Conf. on Software Engineering* (10/76)
T. A. Dolotta *et al.* (p. 25)
Reprinted from that conference's proceedings. Four papers that describe various aspects of PWB/UNIX.

B. Basic User Information:

- B.0 *PWB/UNIX User's Manual—Edition 1.0* (5/77)
T. A. Dolotta, R. C. Haight, and E. M. Piskorik, eds.
Describes all commands, subroutines, and system calls. Furnished as a separate volume. Available on-line.
- B.1 *PWB/UNIX Documentation Roadmap* (5/77)
J. R. Mashey (p. 7)
A terse, up-to-date outline of useful documents and information sources. Available on-line.
- B.2 *PWB/UNIX Beginner's Course* (12/77)
M. E. Pearlman and S. H. Strauss
An outline with view graphs.
- B.3 *A Tutorial Introduction to the UNIX Text Editor* (10/74)
B. W. Kernighan (p. 10)
Expectedly elementary, but very useful for beginners.
- B.4 *Advanced Editing on UNIX* (8/76)
B. W. Kernighan (p. 16)
Meant to help secretaries, typists, and programmers make effective use of UNIX facilities for preparing and editing text.
- B.5 *PWB/UNIX Shell Tutorial* (9/77)
J. R. Mashey (p. 25)
Describes the PWB/UNIX command interpreter.
- B.6 *UNIX for Beginners* (10/74)
B. W. Kernighan (p. 14)
A slightly dated tutorial.
- B.7 *UNIX Programming* (10/75)
B. W. Kernighan and D. M. Ritchie (p. 17)
Introduction to programming on UNIX. The emphasis is on how to write programs that interface with the operating system. Does *not* cover material in *A New Input/Output Package* (item B.10 below).

- B.8 *C Reference Manual* (5/77)
D. M. Ritchie (p. 32)
Terse, but complete.
- B.9 *Programming in C—A Tutorial* (5/75)
B. W. Kernighan (p. 27)
Should be read before tackling the *C Reference Manual* (item B.8 above).
- B.10 *A New Input-Output Package* (7/77)
D. M. Ritchie (p. 6)
Should be used for all new C programs.
- B.11 *A General-Purpose Subroutine Library for PWB/UNIX* (7/77)
A. L. Glasser (p. 7)
Complements *A New Input/Output Package* (item B.10 above).
- B.12 *Guide to IBM Remote Job Entry for PWB/UNIX Users* (9/77)
A. L. Sabsevitz (p. 7)
Describes the RJE facility between a PWB/UNIX system and an IBM System/370.
- B.13 *SCCS/PWB User's Manual* (11/77)
L. E. Bonanni and A. L. Glasser (p. 22)
Describes the Programmer's Workbench Source Code Control System.

T. Text Processing, Formatting, and Typesetting:

- T.1 *NROFF/TROFF User's Manual* (5/77)
J. F. Ossanna (p. 34)
NROFF and TROFF are text processors. NROFF formats text for a variety of typewriter-like terminals. TROFF formats text for a Graphic Systems, Inc. phototypesetter.
- T.2 *PWB/IMM—Programmer's Workbench Memorandum Macros* (10/77)
D. W. Smith and J. R. Mashey (p. 56)
User's guide and reference manual for PWB/IMM, a general-purpose package of text formatting macros for use with NROFF and TROFF.
- T.3 *Typing Documents with PWB/IMM* (10/77)
D. W. Smith and E. M. Piskorik (p. 16)
A fanfold card that fits into a pocket(book).
- T.4 *PWB/IMM Tutorial* (12/77)
N. W. Smith
Introduction to PWB/UNIX text processing.
- T.5 *Tbl—A Program to Format Tables* (9/77)
M. E. Lesk (p. 17)
Preprocessor for TROFF or NROFF that makes even very complex tables easy to specify.
- T.6 *A TROFF Tutorial* (8/76)
B. W. Kernighan (p. 13)
Introduction to the most basic use of TROFF (and, by implication, NROFF).
- T.7 *Typesetting Mathematics—User's Guide (Second Edition)* (6/76)
B. W. Kernighan and L. L. Cherry (p. 11)
Describes the EQN and NEQN preprocessors for TROFF and NROFF, respectively. They allow one to typeset complex formulae, equations, arrays, etc., both in-line and displayed.
- T.8 *New Graphic Symbols for EQN and NEQN* (9/76)
C. Scrocca (p. 8)
Defines a set of special characters frequently used in technical documents. Shows how to use them and discusses what is involved in making a special character in NROFF and TROFF.

- T.9 *PWBIUNIX View Graph and Slide Macros* (12/77)
T. A. Dolotta and D. W. Smith
Greatly eases the task of making transparencies with TROFF.

A. Additional Facilities:

- A.1 *BC—An Arbitrary Precision Desk Calculator Language* (5/75)
L. L. Cherry and R. Morris (p. 14)
A language and a compiler for doing arbitrary-precision arithmetic.
- A.2 *DC—An Interactive Desk Calculator* (5/75)
R. Morris and L. L. Cherry (p. 8)
Interactive desk calculator program that does arbitrary-precision integer arithmetic.
- A.3 *YACC—Yet Another Compiler Compiler* (5/75)
S. C. Johnson (p. 30)
Generates parsers from context-free language specifications.
- A.4 *LEX—Lexical Analyzer Generator* (4/77)
M. E. Lesk and E. Schmidt (p. 13)
LEX helps write programs whose control flow is directed by instances of regular expressions in the input stream.
- A.5 *RATFOR—A Preprocessor for a Rational Fortran* (1/77)
B. W. Kernighan (p. 12)
IF-ELSE, WHILE, and other useful control structures.
- A.6 *The M4 Macro Processor* (4/77)
B. W. Kernighan and D. M. Ritchie (p. 6)
A general-purpose macro language; can be used as a preprocessor for RATFOR, C, etc.
- A.7 *Make—A Program for Maintaining Computer Programs* (4/77)
S. I. Feldman (p. 9)
Make provides a simple mechanism for maintaining up-to-date versions of programs that result from many operations on a number of files.


I. Internals, Operations, and Administration:


- I.1 *Setting Up PWBIUNIX* (9/77)
R. C. Haight, W. D. Roome, and L. A. Wehr (p. 16)
Procedures used to install PWB/UNIX on the PDP-11/45 or /70 and the steps necessary to regenerate all of the PWB/UNIX programs.
- I.2 *Administrative Advice for PWBIUNIX* (10/77)
R. C. Haight (p. 8)
Hints for approaching operational serenity.
- I.3 *PWBIUNIX Operations Manual* (9/77)
M. E. Pearlman (p. 36)
Describes the daily routine at the console. Text (but not pictures) available on-line.
- I.4 *Repairing Damaged PWBIUNIX File Systems* (11/77)
P. D. Wandzilak
Comes in handy after a power failure, etc.
- I.5 *PWBIUNIX RJE Administrator's Guide* (12/77)
A. L. Sabsevitz
What to do when it breaks.

- I.6 *The UNIX I/O System* (6/74)
D. M. Ritchie (p. 9)
Describes how to write device drivers for UNIX.
- I.7 *On the Security of UNIX* (6/74)
D. M. Ritchie (p. 4)
A short, but enlightening, discussion.
- I.8 *UNIX Assembler Reference Manual* (6/73)
D. M. Ritchie (p. 12)
As a last resort ...
- I.9 *PWB/UNIX Manual Page Macros* (8/77)
E. M. Piskorik (p. 7)
Tells how to make PWB/UNIX User's Manual pages.

R. Recommended Reading—not Included:

- R.1 *Software Tools*
B. W. Kernighan and P. J. Plauger (p. 338)
Addison-Wesley, Reading, MA; 1976.
- R.2 The UNIX Command Language
K. Thompson
In *Structured Programming—Infotech State of the Art Report*. Infotech International Limited,
Nicholson House, Maidenhead, Berkshire, England; 1976; pp. 375-84.







**PWB/UNIX—Overview
and
Synopsis of Facilities**

*T. A. Dolotta
R. C. Haight*

June 1977

PWB/UNIX
Overview and Synopsis of Facilities

CONTENTS

OVERVIEW	1
1. UNIX TIME-SHARING SYSTEM	1
2. PWB/UNIX COMMAND LANGUAGE	2
3. DOCUMENT PREPARATION AND TEXT PROCESSING	2
4. REMOTE JOB ENTRY	2
5. SOURCE CODE CONTROL SYSTEM	2
HARDWARE	3
SOFTWARE, FACILITIES, AND DOCUMENTATION	3
1. BASIC SOFTWARE	4
1.1. Operating System 4	
1.2. Access Control 4	
1.3. Manipulation of Files and Directories 4	
1.4. Execution of Programs 6	
1.5. Status Inquiries 7	
1.6. Inter-User Communication 8	
1.7. Program Development Package 8	
1.8. Utilities 9	
1.9. Operations, Maintenance, Backup, and Accounting 10	
2. PROGRAMMING LANGUAGES	11
2.1. The Programming Language C 11	
2.2. FORTRAN 12	
2.3. Other Algorithmic Languages 12	
2.4. Macroprocessors and Compiler-Compilers 13	
3. TEXT PROCESSING	14
3.1. Formatters 14	
3.2. Other Text Processing Tools 15	
4. REMOTE JOB ENTRY TO IBM SYSTEM/370	16
5. TYPESETTING	16
6. DEMONSTRATION AND TRAINING PROGRAMS	17
7. PWB/UNIX USER'S MANUAL	17
8. DOCUMENTS FOR PWB/UNIX	18

PWB/UNIX—Overview and Synopsis of Facilities

T. A. Dolotta

R. C. Haight

Bell Laboratories
Piscataway, New Jersey 08854

OVERVIEW

The Programmer's Workbench (PWB) is a computing facility that provides a convenient working environment and a uniform set of tools for computer program development, as well as for text processing. PWB was developed at Bell Laboratories. As of mid-1977, it supports in excess of 1,000 users. It is based on Bell Laboratories' UNIX* time-sharing system, so that it is more properly known as PWB/UNIX.

1. UNIX TIME-SHARING SYSTEM

The UNIX time-sharing system is a general-purpose, multi-user, interactive operating system specifically engineered to make the designer's, programmer's, and documenter's computing environment simple, efficient, flexible, and productive. PWB/UNIX contains features such as:

- A hierarchical file system.
- A flexible, easy-to-use command language; can be "tailored" to meet specific user needs.
- Ability to execute sequential, asynchronous, and background processes.
- A powerful context editor.
- Very flexible document preparation and text processing systems.
- Access to all the facilities of other ("host" or "target") computer systems, such as the IBM System/370.
- A high-level programming language conducive to structured programming (C).
- The programming languages Basic and FORTRAN.
- Symbolic debugging systems.
- A variety of system programming tools (e.g., compiler-compilers).
- Sophisticated "desk-calculator" packages.

The central processing units (CPUS) for PWB/UNIX are the Digital Equipment Corporation's (DEC†) PDP-11/45 and /70 computers. These computers, because of their price and word size (16 bits per word), are really large minicomputers. Although PWB/UNIX can run on hardware costing as little as \$60,000, a typical PWB/UNIX system costs about \$120,000 and can support 24 simultaneous users with ease. Larger systems can support twice that number. The cost per user-hour of PWB/UNIX is significantly lower than that of most other interactive computer systems. PWB/UNIX typically runs unattended.

The PWB/UNIX file system consists of a highly-uniform set of directories and files arranged in a tree-like hierarchical structure. Some of its features are:

- Simple and consistent naming conventions; names can be absolute, or relative to any directory in the file system hierarchy.
- Mountable and de-mountable file systems and volumes.
- File linking across directories.
- Automatic file space allocation and de-allocation that is invisible to users.
- A complete set of flexible directory and file protection modes; allows all combinations of *read*, *write*, and *execute* access, independently for the owner of each file or directory, for a group of users (e.g., all members of a project), and for all other users; protection modes can be set dynamically.

* UNIX is a Trademark/Service Mark of the Bell System.

† DEC and PDP are registered Trademarks of the Digital Equipment Corporation, Maynard, MA.

- Facilities for creating, accessing, moving, and processing files, directories, or sets of these in a simple, uniform, and natural way.
- Each physical I/O device, from interactive terminals to main memory, is treated like a file, allowing uniform file and device I/O.

2. PWB/UNIX COMMAND LANGUAGE

The PWB/UNIX command language utilizes an extended version of the UNIX *Shell* (command language interpreter), as well as commands designed mainly for use within *Shell procedures* (command files). These PWB/UNIX extensions to the UNIX Shell have been aimed at improving its use by large programming groups, and at making it even more convenient for use as a high-level programming language. In line with the philosophy of much existing UNIX software, a very strong attempt has been made to add new features only when they were shown to be necessary by actual user experience, in order to avoid contaminating a compact, elegant system with unnecessary frills that would destroy the inherent simplicity of the design of UNIX. By utilizing the Shell as a programming language, PWB/UNIX users have been able to eliminate a great deal of the programming drudgery that often accompanies a large project. Many manual procedures have been quickly, inexpensively, and conveniently automated. Because it is so easy to create and use Shell procedures, each project that uses PWB/UNIX has tended to customize the general PWB/UNIX environment into one tailored to its own requirements, organizational structure, and terminology.

3. DOCUMENT PREPARATION AND TEXT PROCESSING

In a software development project of any appreciable size, the production of usable, accurate documentation may well consume more effort than the production of the software itself. Several years of experience with many projects that use PWB/UNIX have shown that document preparation should not be separated from software development, and that the combination of a flexible operating system, a powerful command language, and good text processing facilities permits quick and convenient production of many kinds of documentation that might be otherwise unobtainable, impractical, or very expensive.

In PWB/UNIX, one also obtains a very useful "word processing" system—an editing system, text formatting systems, a typesetting system, and spelling and typographical error-detection facilities. The document preparation and text processing facilities of PWB/UNIX include commands that automatically control pagination, style of paragraphs, line justification, hyphenation, multi-column pages, footnote placement, generation of marginal revision bars, generation of tables of contents, etc., for specialized documents such as program run books, or for general documents such as letters, memoranda, legal briefs, etc. There are also excellent facilities for formatting and typesetting complex tables and equations. This document was produced in its entirety by these facilities.

4. REMOTE JOB ENTRY

The PWB/UNIX RJE facility provides for the submission and retrieval of jobs from an IBM host system (e.g., a System/360 or System/370 computer using HASP, ASP, or JES2). To the host system, RJE appears to be a card reader/punch and line-printer station. At the request of a PWB/UNIX user, RJE gathers the job control statements to be sent to the host system and, subsequently, retrieves from the host the resulting output, places it in a convenient PWB/UNIX file for later perusal, and notifies the user of the output's arrival.

5. SOURCE CODE CONTROL SYSTEM

The PWB/UNIX Source Code Control System (SCCS) is an integrated set of commands designed to help software development projects control changes to source code and to files of text (e.g., manuals). It provides facilities for storing, updating, and retrieving, by version number or date, all versions of source code modules or of documents, and for recording who made each software change, when it was made, and why. SCCS is designed to solve most of the source code and documentation control problems that software development projects encounter when customer support, system testing, and development are all proceeding simultaneously.

Some of the main characteristics of SCCS are:

- The exact source code or text, as it existed at any point of development or maintenance, can be recreated at any later time.
- All releases and versions of a source code module or document are stored together, so that common code or text is stored only once.
- Releases in production or system-test status can be protected from unauthorized changes.
- Enough identifying information can be automatically inserted into source code modules to enable one to identify the exact version and release of any such module, given only the corresponding load module or its memory dump.

HARDWARE

PWB/UNIX runs on a DEC PDP-11/45 or /70 with at least the following equipment:

- 96K words of memory (2 bytes/word), with memory management.
- Disks: RP03, RP04, RP05, RP06, or equivalent.
- Console typewriter terminal.
- Clock: KW11L or KW11P.
- Tape: TU10, TU16, or equivalent.
- Floating point: FP11B or FP11C.

The following equipment is strongly recommended:

- DH11 communications controller(s) with full modem control.
- Full-duplex 96-character ASCII terminals.
- Extra disk drive for system backup.
- DQS11B communications controller(s) for RJE.

The minimum memory and disk space specified is enough to run and maintain PWB/UNIX. More will be needed to keep all source on line, or to handle a large number of users, big data bases, diversified complements of devices, or large programs. PWB/UNIX does swapping and sharing of reentrant user code in order to minimize main memory requirements. The resident PWB/UNIX operating system uses 40-48K words, depending on the configuration.

A large PWB/UNIX configuration (PDP-11/70, 256K words of main memory, fixed and moving head disks) can generally provide reasonable service to between 40 and 48 simultaneous users.

SOFTWARE, FACILITIES, AND DOCUMENTATION

The most-often-used PWB/UNIX commands are listed below. Every command, including all options, is issued as a single line, unless specifically described below as being "interactive." Interactive programs can be made to run from a prepared "script" simply by redirecting their input. All commands are described in the PWB/UNIX User's Manual (see Section 7 below). Commands for which additional manuals and tutorials are provided are marked with [m] and [t], respectively. All manuals and tutorials are listed in Section 8 below.

File processing commands that go from standard input to standard output are called "filters" and are marked with [f]. The "pipe" facility of the Shell may be used to connect filters directly to the input or output of other filters and programs.

Almost all of PWB/UNIX is written in C. PWB/UNIX is totally self-supporting: it contains all the software that is needed to generate it, maintain it, and modify it. Source code is included except as noted below.

The system is distributed on three 9-track, 800 BPI magnetic tapes. The programs listed in Section 5 below (*Typesetting*) are distributed on a separate tape at an additional cost.

1. BASIC SOFTWARE

It includes the time-sharing operating system with utilities, an assembler, and a compiler for the programming language C—enough software to regenerate, maintain, and modify PWB/UNIX itself, and to write and run new applications.

1.1. Operating System

- UNIX [m,t] The basic resident code on which everything else depends. Executes the system calls, maintains the file system, and manages the system's resources; contains device drivers, I/O buffers, and other system information. A general description of UNIX design philosophy and system facilities appeared in the *Communications of the ACM* (see Section 8 below). Further capabilities include:
 - Automatically-supported reentrant code.
 - Separation of instruction and data spaces.
 - Timer-interrupt sampling and interprocess monitoring for debugging and measurement.
- Devices All I/O is logically synchronous. Normally, automatic buffering by the system makes the physical record structure invisible and exploits the hardware's ability to do overlapped I/O. Unbuffered physical record I/O is available for unusual applications. Drivers for the following devices are available (others can be written easily):
 - Asynchronous interfaces: DH11, DL11; provide support for the most common interactive ASCII terminals.
 - Synchronous interfaces: DU11, DQS11B.
 - Automatic calling unit interface: DN11.
 - Line printer: LP11, or equivalent.
 - Magnetic tape drives: TU10 and TU16, or equivalent.
 - Fixed head disks: RS03 and RS04, or equivalent.
 - Removable-pack disks: RP03 (recommended only for PDP-11/45) and RP04-5-6, or equivalent, with one or more logical devices per physical device and minimum-latency seek scheduling.
 - Physical memory of PDP-11, or mapped memory in resident system.
 - Phototypesetter: Graphic Systems, Inc., System/1, connected through a DR11C.

1.2. Access Control

- LOGIN Signs on a new user:
 - Adapts to characteristics of terminal.
 - Verifies password and establishes user's individual and group (project) identity.
 - Establishes working directory.
 - Publishes message of the day.
 - Announces presence of mail.
 - Starts command interpreter (*Shell*) or other user-specified program.
- PASSWD Changes a password:
 - User can change own password.
 - Passwords are kept in encrypted form.

1.3. Manipulation of Files and Directories

- CAT [f] Concatenates one or more files onto standard output. Mostly used for unadorned printing, for inserting data into a "pipe," and for buffering output that comes in dribs and drabs.

- PR [f] Prints files with title, date, and page number on every page:
 - Multi-column output.
 - Parallel column merge of several files.
- ED [f,m,t] Interactive context editor. Random access to all lines of a file. It can:
 - Find lines by number or pattern (regular expressions). Patterns can include: specified characters, "don't care" characters, choices among characters, (specified numbers of) repetitions of these constructs, beginning of line, end of line.
 - Add, delete, change, copy, or move lines.
 - Permute contents of a line.
 - Replace one or more instances of a pattern within a line.
 - Combine or split lines.
 - Combine or split files.
 - Escape to Shell (PWB/UNIX command language) during editing.
 - Do any of above operations on every line (in a given range) that matches a pattern.
- SED [f] A stream (one-pass) editor with facilities similar to those of ED.
- REFORM [f] Reformats a file, especially source code from another computer.
- SPLIT Splits a large file into more manageable pieces.
- CSPLIT Like SPLIT, with the splitting controlled by context.
- SUM Computes the check sum of a file; often used to verify that all of a file can be read.
- DD [f] Physical file format translator, for exchanging data with non-UNIX systems, especially OS/360, VS1, MVS, etc.
- CP Copies one file to another. Works on any file regardless of its contents.
- RM Removes file(s). If any names are linked to the file, only the name being removed goes away.
- LN Links another name (alias) to an existing file.
- MV Moves a file. Used for renaming files or directories.
- CHMOD Changes access permissions on file(s). Executable by owner of file(s), or by the super-user.
- CHOWN Changes owner of file(s).
- MKDIR Makes one or more new directories.
- RMDIR Removes one or more (empty) directories.
- CHDIR Changes working (i.e., current) directory.
- FIND Searches the directory hierarchy for, and performs specified commands on, every file that meets given criteria:
 - File name matches a given pattern.
 - Modified date in given range.
 - Date of last use in given range.
 - Given permissions.
 - Given owner.
 - Given special file characteristics.
 - Any logical combination of the above.
 - Any directory can be the starting "node."

- DSW Interactively steps through a directory, deleting or keeping files.
- AR Maintains archives and libraries. Combines several files into one for house-keeping efficiency:
 - Creates new archive.
 - Updates archive by date.
 - Replaces or deletes files.
 - Prints table of contents.
 - Retrieves from archive.
- CPIO Copies a sub-tree of the file system (directories, links, and all) to another place in the file system. Can also copy a sub-tree onto a tape, and later recreate it from tape.
- SCCS [m] SCCS (Source Code Control System) is a collection of PWB/UNIX commands (some interactive) for controlling changes to files of text (typically, the source code of programs or text of documents). It provides facilities for:
 - Storing, updating, and retrieving any version of any source or text file.
 - Controlling updating privileges.
 - Identifying both source and object (or load) modules by version number.
 - Recording who made each change, when it was made, and why.

1.4. Execution of Programs

- SH [f,t] The Shell, or command language interpreter, understands a set of constructs that constitute a full programming language; it allows a user or a command procedure to:
 - Supply arguments to and run any executable program.
 - Very easily redirect standard input and/or standard output.
 - Compose compound commands using the following operators:
 - ; for sequential execution.
 - | for simultaneous execution with output of one process "piped" to the input of another.
 - & for asynchronous operation.
 - () parentheses for grouping.
 - || and && for left-to-right conditional evaluation.
 - Use 26 string variables (*a-z*), and a set of string and integer operators.
 - Control the order in which directories are searched for commands.
 - Trace execution of commands for debugging.
 - Execute Shell procedures, which are command scripts with substitutable arguments.
 - Construct arguments that denote all file names that match a specified pattern (regular expression).
 - Use the Shell itself as a command (recursively).
 - Collect command usage statistics.
 - Tutorial covers most of the commands in this section.
- IF Conditionally executes commands on basis of:
 - String comparison.
 - File existence, type, and accessibility.
 - IF-THEN-ELSE-ENDIF control constructs.
 - Return code of a process.
- WHILE Shell procedure looping; END shows end of loop; BREAK terminates loop; CONTINUE requests next iteration; condition testing as in IF.

- SWITCH Multi-way branch; ENDSW shows end of switch; BREAKSW ends each choice.
- SHIFT Shifts Shell's arguments one argument to the left; allows a Shell procedure to iterate over its arguments.
- = (equals) Assigns values to Shell variables.
- EXPR Evaluates expressions containing Shell variables; knows about arithmetic, logical, and string operators.
- GOTO Goes to a label in a Shell procedure.
- EXIT Terminates a Shell procedure; useful with IF.
- NEXT Names a new standard input for the current Shell procedure.
- PUMP [f] Creates an instance of the current Shell to substitute Shell arguments and variables into the standard input.
- ECHO Prints its arguments on the standard output. Useful for diagnostics or prompts in Shell procedures, or for inserting data into a "pipe."
- ONINTR Traps terminal interrupts.
- SLEEP Suspends execution for a specified time.
- WAIT Waits for termination of processes that are running asynchronously.
- NOHUP Runs a command immune to "hanging up" the terminal.
- NICE Runs a command in low (or high) priority.
- KILL Terminates named process(es).
- CRON Performs actions at specified times:
 - Actions are arbitrary Shell procedures or executable programs.
 - Times are conjunctions of month, day of month, day of week, hour, and minute. Ranges are specifiable for each.
- TEE [f] Passes data between processes (like a "pipe"), but also diverts copies into one or more files.
- RSH Restricted Shell; restricts a user to a subset of PWB/UNIX commands.
- HELP Explains error messages from other programs.

1.5. Status Inquiries

- LS Lists the names of one, several, or all files in one or more directories:
 - Alphabetic or chronological sorting, up or down.
 - Optional information: size, owner, group, date last modified, date last accessed, permissions.
- FILE Tries to determine what kind of information is in a file by consulting the file system index and by reading the file itself.
- DATE Print current date and time. Has considerable knowledge of calendrical and horologic peculiarities; can be used to set PWB/UNIX's idea of date and time. (As yet, cannot cope with Daylight Saving Time in the Southern Hemisphere.)
- DF Reports amount of free space in file system.
- DU Prints a summary of total space occupied by all files in a hierarchy.
- TTY Prints the "name" of your terminal (i.e., the name of the port to which your terminal is connected).

- WHO Tells who is logged into the system:
 - Lists logged-in users, their ports, and time they logged in.
 - Optional history of all logins and logouts.
 - Tells you who you are logged in as.
- PS Reports on active processes:
 - Lists your own or everybody's processes.
 - Tells what commands are being executed at the moment.
 - Optional status information: state and scheduling information, priority, attached terminal, what the process is waiting for, and its size.
- PWD Prints name of your working (i.e., current) directory.
- STTY Sets up options for optimal control of a terminal. In so far as they are deducible from the input, these options are set automatically by LOGIN:
 - Parity.
 - Mapping of upper-case characters to lower case.
 - Carriage-return plus line-feed versus new-line.
 - Interpretation of tab characters.
 - Delays for tab, new-line, and carriage-return characters.
 - Raw versus edited input.
- TABS Sets terminal's tab stops. Knows several "standard" formats: FORTRAN, PL/I, COBOL; eases preparation of input for various RJE host systems.

1.6. Inter-User Communication

- MAIL Mails a message to one or more users. Also used to read and dispose of incoming mail. The presence of mail is announced by LOGIN.
- WRITE Establishes direct, interactive terminal-to-terminal communication with another user.
- WALL Writes to all users who are logged in.
- MMSG Inhibits or permits receipt of messages from WRITE and WALL.

1.7. Program Development Package

A kit of fundamental programming tools. Some of these utilities are used as integral parts of the higher-level languages described in Section 2 below.

- AS [m] Assembler:
 - Creates object program consisting of:
 - Code, possibly read-only.
 - Initialized data or read-write code.
 - Uninitialized data.
 - Relocatable object code is executable without further transformation.
 - Object code normally includes a symbol table (that can be deleted).
 - Several source files can be assembled together.
 - Local labels.
 - Conditional assembly.
 - "Conditional jump" instructions become branches or branches plus jumps, depending on distance.
- Library The basic run-time library. It is used freely by all system software:
 - Number conversions.
 - Time conversions.
 - Mathematical functions: *sin*, *cos*, *log*, *exp*, *atan*, *sqrt*, *gamma*.
 - Buffered character-by-character I/O.

- Random number generator.
 - An elaborate library for formatted I/O [m].
 - Password encryption.
- DB Interactive post-mortem debugger. Works on core dump files produced by program aborts, on object files, and on arbitrary files:
 - Symbolic addressing within files that have symbol tables.
 - Octal, decimal, or ASCII output.
 - Symbolic disassembly.
 - Octal or decimal patching.
 - OD [f] Dumps any file:
 - Output options include: octal or decimal by words, octal by bytes, ASCII, operation codes, hexadecimal, or any combination thereof.
 - Range of dumping is controllable.
 - LD Linkage editor. Combines relocatable object files. Inserts required routines from specified libraries; resulting code:
 - Can be made sharable.
 - Can be made to have separate instruction and data spaces.
 - NM Prints the *namelist* (symbol table) of an object program. Provides control over the style and order of names that are printed.
 - SIZE Reports the main memory requirements of one or more object files.
 - STRIP Removes the relocation and symbol table information from an object file to save file space.
 - PROF Constructs a profile of time spent in each routine from data gathered by time-sampling the execution of a program; gives subroutine call frequencies and average times for C programs.

1.8. Utilities

- CREF Makes cross-reference listings of a set of files. Each symbol is listed together with file name, line number within file, and text of each line in which it occurs:
 - Assembler or C language files can be cross-referenced.
 - Gathers or suppresses references to selected symbols.
 - Various ways to sort output are available.
 - Allows selective printing of uniquely occurring symbols.
- SORT [f] Merges and/or sorts ASCII files line-by-line:
 - In ascending or descending order.
 - Lexicographically or on numeric key.
 - On multiple keys located by delimiters or by position.
 - Can fold upper-case characters together with lower-case into dictionary order.
- UNIQ [f] Deletes successive duplicate lines in a file:
 - Prints lines that were originally unique, duplicated, or both.
 - Can give redundancy count for each line.
- TR [f] Does character translation according to an arbitrary code:
 - Can "squeeze out" repetitions of selected characters.
 - Can delete selected characters.

- DIFF [f] Reports line changes, additions, and deletions necessary to bring two files into agreement; can produce an editor script to convert one file into another.
- COMM [f] Identifies common lines in two sorted files. Output in up to 3 columns shows lines present in first file only, present in second file only, and/or present in both.
- CMP Compares two files and reports disagreeing bytes.
- GREP [f] Prints all lines in one or more files that match a pattern of the kind used by ED (the editor):
 - Can print all lines that fail to match.
 - Can print count of "hits."
- WC [f] Counts lines and "words" (strings separated by blanks or tab characters) in a file.
- GATH [f] Gathers real and virtual files, including the output of other commands; performs keyword substitution, expansion of tabs, nested file inclusion, etc.
- MAKE [m] Provides a general mechanism for keeping up-to-date large programs that are created by performing many operations on a number of files.
- TIME Runs a command and reports timing information about it.

1.9. Operations, Maintenance, Backup, and Accounting

- MOUNT Attaches a device containing a file system to the tree of directories. Protects against nonsense arrangements.
- UMount Removes the file system contained on a device from the tree of directories. Protects against removing a busy device.
- MKFS Makes a new file system on a device.
- MKNOD Makes an i-node (file system entry) for a special file. Special files are physical devices, virtual devices, physical memory, etc.
- TP Manages file archives on magnetic tape:
 - Collects files into an archive.
 - Updates archive by date.
 - Replaces or deletes archived files.
 - Prints table of contents of an archive.
 - Retrieves from an archive.
- DUMP Dumps the file system on the specified device, selectively by date, or indiscriminately.
- RESTOR Restores a dumped file system, or selectively retrieves parts thereof.
- VOLCOPY An alternate file backup arrangement; system should have a spare disk drive. Protective labeling of disks and tapes is included.
- SU Temporarily allows one to become the super-user with all the rights and privileges thereof. Requires a (distinct) password.
- CHECK Together with DCHECK, ICHECK, and NCHECK, is used to check the consistency of file systems and directories; among them, these programs can:
 - Print statistics: number of files, number of directories, number of special files, space used, free space.
 - Report duplicate use of space.
 - Retrieve lost space.

- Report inaccessible files.
 - Check consistency of directories.
 - List names of all files.
 - Reorganize free disk space for maximum operating efficiency.
- CLRM Peremptorily expunges a file and its space from a file system. Used to repair damaged file systems.
- SYNC Forces all outstanding I/O on the system to completion. Used to shut down the system gracefully.
- CONFIG Tailors device-dependent system code to a specific hardware configuration. As distributed, PWB/UNIX can be brought up directly on any acceptable CPU with an acceptable tape drive and disk, sufficient amount of main memory, a console terminal, and either clock.
- AC Publishes cumulative connect-time report:
 - Connect time by user or by day.
 - For all users, or for selected users.
- SA Prints Shell accounting report. For each command executed, gives:
 - Number of times used.
 - Total system time, user time, and elapsed time.
 - Optionally, averages and percentages.
 - Output can be sorted on various fields.
- LASTCOM Searches Shell accounting records in reverse chronological order for commands executed by the given user(s), or from the given terminal(s).
- DCAT Transfers files between two PWB/UNIX systems over a medium-speed, dial-up synchronous link.
- Operations A number of commands and several manuals (see Section 8 below) are provided for the operations staff.

2. PROGRAMMING LANGUAGES

2.1. The Programming Language C

- CC [m,t] Compiles and/or link-edits programs in the C language. The PWB/UNIX operating system, almost all of its subsystems, and C itself are written in C:
 - General-purpose language designed for structured programming.
 - Data types:
 - Character.
 - Integer.
 - Long integer.
 - Floating-point.
 - Double.
 - Pointers to all types.
 - Functions returning all types.
 - Arrays of any type.
 - Structures containing various types.
 - Provides machine-independent control of all machine facilities, including to-memory operations and pointer arithmetic.
 - Macro-preprocessor for parameterized code and for the inclusion of other files.
 - All procedures recursive, with parameters passed by value.
 - Object code uses full addressing capability of the PDP-11.
 - Run-time library gives access to all system facilities.

- CDB An interactive debugger tailored for use with C:
 - Usable in real time or post-mortem.
 - The debugger is a completely separate process from the program being debugged. No debugging code is loaded with the program being debugged.
 - Prints in the most appropriate notation: characters, integers (octal and decimal), floating-point numbers, double precision numbers, machine instructions (disassembled).
 - Stack trace and fault identification.
 - Break-point tracing.
- ADB A more recent, more powerful, but also more complex interactive debugger than CDB.
- CB [f] C beautifier: gives a C program that well-groomed, structured, indented look.

2.2. FORTRAN

- FC Compiles and/or link-edits FORTRAN IV programs. Object code is partially interpreted.
 - Idiosyncrasies:
 - Free-form, lower-case source code.
 - No arithmetic statement functions.
 - Unformatted I/O requires that expected and actual record lengths agree.
 - No BACKSPACE statement.
 - No P format control character on input.
 - Handles mixed-mode arithmetic, general subscripts, and general DO limits.
 - 32-bit integer arithmetic.
 - Free-format numeric input.
 - Understands these nonstandard specifications:
 - LOGICAL*1, *2, *4
 - INTEGER*2, *4
 - REAL*4, *8
 - COMPLEX*8, *16
 - IMPLICIT
- RC [m] Compiles and/or link-edits RATFOR programs. RATFOR adds rational control structures to FORTRAN:
 - ELSE, FOR, WHILE, REPEAT ... UNTIL statements.
 - Symbolic constants.
 - File inclusion.
 - Compound statements.
 - Can produce genuine FORTRAN to carry away.

2.3. Other Algorithmic Languages

- BAS An interactive interpreter, similar in style to Basic. Interprets unnumbered statements immediately, numbered statements when given the "run" command:
 - Statements include:
 - comment
 - dump
 - for ... next
 - goto

- if ... else ... fi
 - list
 - print
 - prompt
 - return
 - run
 - save
- All calculations in double precision.
- Recursive function defining and calling.
- Built-in functions include *log, exp, sin, cos, atn, int, sqr, abs, rnd*.
- Escape to the standard PWB/UNIX editor for complex program editing.
- DC [m] Interactive programmable desk calculator. Has named storage locations, as well as conventional stack for holding integers and programs:
 - Arbitrary-precision decimal arithmetic.
 - Appropriate treatment of decimal fractions.
 - Arbitrary input and output radices, in particular binary, octal, decimal, and hexadecimal.
 - Postfix (“Reverse Polish”) operators:
 - + - * /
 - remainder, power, square root
 - load, store, duplicate, clear
 - print, enter program text, execute
- BC [m] A C-like interactive interface to the desk calculator DC:
 - All the capabilities of DC with a high-level syntax.
 - Arrays and recursive functions.
 - Immediate evaluation of expressions and evaluation of functions upon call.
 - Arbitrary-precision elementary functions: *exp, sin, cos, atan*.
 - Goto-less programming.
- SNO An interpreter very similar to SNOBOL 3; its limitations are:
 - Function definitions are static.
 - Pattern matches are always anchored.
 - No built-in functions.

2.4. Macroprocessors and Compiler-Compilers

- M4 [f,m] A general-purpose macroprocessor:
 - Stream-oriented, recognizes macros anywhere in text.
 - Integer arithmetic.
 - String and substring capabilities.
 - Condition testing, file manipulation, arguments.
- YACC [m] An LR(1)-based compiler-writing system. During execution of resulting parsers, arbitrary C functions can be called to do code generation or take semantic actions:
 - BNF syntax specifications.
 - Precedence relations.
 - Accepts formally ambiguous grammars with non-BNF resolution rules.
- LEX [m] LEX helps write programs whose control flow is directed by instances of regular expressions in the input stream. It is well suited for editor-script type transformations and for segmenting input in preparation for a parsing routine.

3. TEXT PROCESSING

3.1. Formatters

High-level formatting macros have been developed to ease the task of preparing documents with NROFF, as well as to exploit its more complex formatting capabilities.

- NROFF [f,m,t] Advanced formatter for terminals. Capable of many elaborate feats:
 - Justification of either or both margins.
 - Automatic hyphenation.
 - Generalized page headers and footers, automatic page numbering, with even-odd page differentiation capability, etc.
 - Hanging indents and one-line indents.
 - Absolute and relative parameter settings.
 - Optional legal-style numbering of output lines.
 - Nested or chained input files.
 - Complete page format control, keyed to dynamically-planted "traps" at specified lines.
 - Several separately-definable formatting environments (e.g., one for regular text, one for footnotes, and one for "floating" tables and displays).
 - Macros with substitutable arguments.
 - Conditional execution of macros.
 - Conditional insertion or deletion of text.
 - String variables that can be invoked in mid-line.
 - Computation and printing of numerical quantities.
 - String-width computations for unusually-difficult layout problems.
 - Positions and distances expressible in inches, centimeters, ems, ens, line spaces, points, picas, machine units, and arithmetic combinations thereof.
 - Dynamic (relative or absolute) positioning.
 - Horizontal and vertical line drawing.
 - Multi-column output on terminals capable of reverse line-feed, or through the postprocessor COL.

- -MM [m,t] A standardized manuscript layout for use with NROFF/TROFF. Provides a flexible, user-oriented interface to these two formatters; designed to be:
 - Robust in face of user errors.
 - Easy to use by all levels of users.
 - Adaptable to a wide range of output styles.
 - Can be extended by users familiar with the formatter.
 - Compatible with both NROFF and TROFF.

Some of its features are:

- Page numbers and draft dates.
 - Cover sheets and title pages.
 - Automatically-numbered or "lettered" headings.
 - Automatically-numbered or "lettered" lists.
 - Automatically-numbered figure and table captions.
 - Automatically-numbered and positioned footnotes.
 - Single- or double-column text.
 - Paragraphing, displays, and indentation.
 - Automatic table of contents.
- NEQN [f,m] A mathematical preprocessor for NROFF with the same facilities as EQN (see Section 5 below), except for the limitations imposed by the graphic capabilities of the terminal being used. Prepares formulae for display on the Model 37 TELETYPE[®], various Diablo-mechanism terminals, etc.

- TBL [f,m] A preprocessor for NROFF that translates simple descriptions of table layouts and contents into detailed formatting instructions:
 - Computes appropriate column widths.
 - Handles left- and right-justified columns, centered columns, and decimal-point aligned columns.
 - Places column titles; spans these titles, as appropriate.

For example:

Composition of Foods			
<i>Food</i>	<i>Percent by Weight</i>		
	Protein	Fat	Carbo- hydrate
Apples	.4	.5	13.0
Halibut	18.4	5.2	...
Lima beans	7.5	.8	22.0
Milk	3.3	4.0	5.0
Mushrooms	3.5	.4	6.0
Rye bread	9.0	.6	52.7

3.2. Other Text Processing Tools

- SPELL Finds spelling errors by looking up all uncommon words from a document in a large spelling list. Knows about prefixes and suffixes and can cope with such rotten spellings as "roted."
- TYPO Finds typographical errors. Statistically analyzes all the words in a text, weeds out several thousand familiar ones, and prints the rest sorted so that the most improbably spelled ones tend to come to the top of the list.
- DIFFMARK [f] Judiciously combines DIFF, ED, and NROFF/TROFF to produce marginal revision bars on output.
- PTX Generates a permuted index, like the one in the PWB/UNIX User's Manual.
- GRAPH [f] Given the coordinates of the points to be plotted, draws the corresponding graph; has many options for scaling, axes, grids, labeling, etc.
- PLOT [f] Makes the output of GRAPH suitable for plotting on a Diablo-mechanism terminal.
- GSI [f] Exploits the hardware facilities of GSI300, DASI300, and other Diablo-mechanism terminals:
 - Implements reverse line-feeds and forward and reverse fractional-line motions.
 - Allows any combination of 10- or 12-pitch printing with 6 or 8 lines/inch spacing.
 - Approximates Greek letters and other special characters by overstriking in plot mode.
- 450 [f] Like GSI, but for the newer Diablo-mechanism terminals (e.g., DASI450).
- HP [f] Like GSI, but for the Hewlett-Packard 2640 family of terminals.
- COL [f] Reformats files with reverse line-feeds so that they can be correctly printed on terminals that cannot reverse line-feed.

4. REMOTE JOB ENTRY TO IBM SYSTEM/370

- SEND [m] User interface to PWB/UNIX Remote Job Entry (RJE) facility; may be used interactively. RJE provides for submission and retrieval of jobs from a host system (e.g., an IBM System/360 or System/370 computer using HASP, ASP, or JES2). To the host system, RJE appears to be an IBM 2770 card reader/punch and line-printer station. At the request of a PWB/UNIX user, RJE gathers job control and source statements to be sent to the host system, and subsequently returns to the user output listings and object "decks" (in the form of PWB/UNIX files).
- RJESTAT Reports interactively on the status of any job(s) on the RJE host systems, as well as on the status of the RJE links themselves.

5. TYPESETTING

This software generates output on a Graphic Systems, Inc., C/A/T System/1 phototypesetter. It is distributed separately as an (extra cost) add-on to PWB/UNIX.

TROFF is much like NROFF (see Section 3.1 above), but provides a number of different character styles, a selection of character (point) sizes, and proportional spacing; -MM, TBL, and the typesetting language for mathematics (EQN/NEQN) are compatible with both NROFF and TROFF, so that the user-level formatting instructions can be the same for both.

- TROFF [f,m,t] Provides facilities that are upward-compatible with NROFF, but with the following additions:
 - Vocabulary of several 102-character fonts (any 4 simultaneously) in 15 different point sizes.
 - Character-width and string-width computations for unusually difficult layout problems.
 - Overstrikes and built-up brackets.
 - Dynamic (relative or absolute) point size selection, globally or at the character level.
 - Terminal output for rough sampling of the product.

■ This entire document was typeset by TROFF, assisted by -MM, TBL, and EQN. ■

- EQN [f,m] A mathematical preprocessor for TROFF. Translates in-line or displayed formulae from a very easy-to-type form into detailed typesetting instructions. For example:

sigma sup 2 = 1 over N sum from j=1 to N (x sub j - x bar) sup 2

produces:

$$\sigma^2 = \frac{1}{N} \sum_{j=1}^N (x_j - \bar{x})^2$$

- Automatic calculation of point size changes for subscripts, superscripts, sub-subscripts, etc.
 - Full vocabulary of Greek letters, such as γ , Π , Γ , α .
 - Automatic calculation of the size of large brackets.
 - Vertical "piling" of formulae for matrices, conditional alternatives, etc.
 - Integrals, sums, etc., with arbitrarily complex limits.
 - Diacriticals: dots, double dots, hats, bars, etc.
 - Easily learned by non-programmers and mathematical typists.
 - Formulae can appear within tables to be formatted by TBL (see Section 3.1 above).
- -MV [m] A set of easy-to-use TROFF macros for making view graphs and slides. Intended for use by secretaries and others not trained in "graphic arts." All facilities of TROFF, TBL, and EQN are available.

6. DEMONSTRATION AND TRAINING PROGRAMS

Unless otherwise indicated below, source code for the following interactive programs is *not* included:

- QUIZ Tests your knowledge of Shakespeare, presidents, capitals, etc. Source code included.
- CHESS This chess-playing program scored 1-2-1 and 3-0-1 in the 1973 and 1974 ACM Computer Chess Championships.
- BJ A blackjack dealer.
- REVERSI A game of dramatic reversals.
- MOO A fascinating number-guessing game, rather like Mastermind®.
- FACTOR Computes prime factors of a number.
- CAL Prints a calendar of specified month or year between A.D. 1 and 9999. Source code included.
- UNITS Converts quantities between different scales of measurement. Knows hundreds of units; for example, how many kilometers/second (or furlongs/fortnight) is a parsec/megayear? Source code included.
- TTT A (traditional, 3×3, two-dimensional) tic-tac-toe program that learns. It never makes the same mistake twice, unless you make it forget what it has learned.
- CUBIC An accomplished player of 4×4×4 tic-tac-toe.
- WUMP Thrilling hunt for the mighty wumpus in a dangerous cave.

7. PWB/UNIX USER'S MANUAL

- MAN Prints one or more specified PWB/UNIX User's Manual entries on your terminal. Machine-readable and hard-copy versions of that manual are provided:
 - System overview.
 - All commands.
 - All system calls.
 - Most subroutines in the assembler, C, and FORTRAN libraries.
 - All devices and other special files.
 - File formats for all files known to the system software.
 - Descriptions of various terminals commonly used with PWB/UNIX.
 - "Boot" procedures.

8. DOCUMENTS FOR PWB/UNIX

In addition to the PWB/UNIX User's Manual—Edition 1.0, which is a separate volume (see Section 7 above), the following are provided with PWB/UNIX:

General:

1. PWB/UNIX—Overview and Synopsis of Facilities (this document).
2. The UNIX Time-Sharing System (reprinted from *Comm. ACM* 17(7):365-75, July 1974).
3. The UNIX Time-sharing System—A Retrospective.
4. PWB Papers from the Second International Conference on Software Engineering.

Basic User Information:

1. PWB/UNIX Documentation Roadmap.
2. PWB/UNIX Beginner's Course.
3. A Tutorial Introduction to the UNIX Text Editor.
4. Advanced Editing on UNIX.
5. PWB/UNIX Shell Tutorial.
6. UNIX for Beginners.
7. UNIX Programming.
8. C Reference Manual.
9. Programming in C—A Tutorial.
10. A New Input-Output Package.
11. A General-Purpose Subroutine Library for PWB/UNIX.
12. Guide to IBM Remote Job Entry for PWB/UNIX Users.
13. SCCS/PWB User's Manual.

Text Processing, Formatting, and Typesetting:


1. NROFF/TROFF User's Manual.
2. PWB/MM—Programmer's Workbench Memorandum Macros.
3. Typing Documents with PWB/MM.
4. PWB/MM Tutorial.
5. Tbl—A Program to Format Tables.
6. A TROFF Tutorial.
7. Typesetting Mathematics—User's Guide (Second Edition).
8. New Graphic Symbols for EQN and NEQN.
9. PWB/UNIX View Graph and Slide Macros.


Additional Facilities:

1. BC—An Arbitrary Precision Desk Calculator Language.
2. DC—An Interactive Desk Calculator.
3. YACC—Yet Another Compiler Compiler.
4. LEX—Lexical Analyzer Generator.
5. RATFOR—A Preprocessor for a Rational FORTRAN.
6. The M4 Macro Processor.
7. Make—A Program for Maintaining Computer Programs.

Internals, Operations, and Administration:

1. Setting up PWB/UNIX.
2. Administrative Advice for PWB/UNIX.
3. PWB/UNIX Operations Manual.
4. Repairing Damaged PWB/UNIX File Systems.
5. PWB/UNIX RJE Administrator's Guide.
6. The UNIX I/O System.
7. On the Security of UNIX.
8. UNIX Assembler Reference Manual.
9. PWB/UNIX Manual Page Macros.





The UNIX Time-Sharing System

*Dennis M. Ritchie
Ken Thompson*

*Bell Laboratories
Murray Hill, N. J. 07974*

ABSTRACT

UNIX is a general-purpose, multi-user, interactive operating system for the Digital Equipment Corporation PDP-11/40, 11/45 and 11/70 computers. It offers a number of features seldom found even in larger operating systems, including

1. A hierarchical file system incorporating demountable volumes,
2. Compatible file, device, and inter-process I/O,
3. The ability to initiate asynchronous processes,
4. System command language selectable on a per-user basis,
5. Over 100 subsystems including a dozen languages.

This paper discusses the nature and implementation of the file system and of the user command interface.

1. Introduction

There have been three versions of UNIX. The earliest version (circa 1969-70) ran on the Digital Equipment Corporation PDP-7 and -9 computers. The second version ran on the unprotected PDP-11/20 computer. This paper describes only the PDP-11/40, /45 and /70¹ system, since it is more modern and many of the differences between it and older UNIX systems result from redesign of features found to be deficient or lacking.

Since PDP-11 UNIX became operational in February, 1971, about 100 installations have been put into service; they are generally smaller than the system described here. Most of them are engaged in applications such as the preparation and formatting of patent applications and other textual material, the collection and processing of trouble data from various switching machines within the Bell System, and recording and checking telephone service orders. Our own installation is used mainly for research in operating systems, languages, computer networks, and other topics in computer science, and also for document preparation.

Copyright © 1974, Association for Computing Machinery, Inc. General permission to republish, but not for profit, all or part of this material is granted provided that ACM's copyright notice is given and that reference is made to the publication, to its date of issue, and to the fact that reprinting privileges were granted by permission of the Association for Computing Machinery.

This is a revised version of an article appearing in the Communications of the ACM, Volume 17, Number 7 (July 1974) pp 365-375. That article is a revised version of a paper presented at the Fourth ACM Symposium on Operating Systems Principles, IBM Thomas J. Watson Research Center, Yorktown Heights, New York, October 15-17, 1973.

Perhaps the most important achievement of UNIX is to demonstrate that a powerful operating system for interactive use need not be expensive either in equipment or in human effort: UNIX can run on hardware costing as little as \$40,000, and less than two man-years were spent on the main system software. Yet UNIX contains a number of features seldom offered even in much larger systems. Hopefully, however, the users of UNIX will find that the most important characteristics of the system are its simplicity, elegance, and ease of use.

Besides the system proper, the major programs available under UNIX are

- assembler,
- text editor based on QED²,
- linking loader,
- symbolic debugger,
- compiler for a language resembling BCPL³ with types and structures (C),
- interpreter for a dialect of BASIC,
- phototypesetting and equation setting programs
- Fortran compiler,
- Snobol interpreter,
- top-down compiler-compiler (TMG⁴),
- bottom-up compiler-compiler (YACC),
- form letter generator,
- macro processor (M6⁵),
- permuted index program.

There is also a host of maintenance, utility, recreation and novelty programs. All of these programs were written locally. It is worth noting that the system is totally self-supporting. All UNIX software is maintained under UNIX; likewise, this paper and all other UNIX documents were generated and formatted by the UNIX editor and text formatting program.

2. Hardware and software environment

The PDP-11/45 on which our UNIX installation is implemented is a 16-bit word (8-bit byte) computer with 112K bytes of core memory; UNIX occupies 53K bytes. This system, however, includes a very large number of device drivers and enjoys a generous allotment of space for I/O buffers and system tables; a minimal system capable of running the software mentioned above can require as little as 64K bytes of core altogether.

Our PDP-11 has a 1M byte fixed-head disk, used for file system storage and swapping, four moving-head disk drives which each provide 2.5M bytes on removable disk cartridges, and a single moving-head disk drive which uses removable 40M byte disk packs. There are also a high-speed paper tape reader-punch, nine-track magnetic tape, and DECTape (a variety of magnetic tape facility in which individual records may be addressed and rewritten). Besides the console typewriter, there are 30 variable-speed communications interfaces attached to 100-series datasets and a 201 dataset interface used primarily for spooling printout to a communal line printer. There are also several one-of-a-kind devices including a Picturephone[®] interface, a voice response unit, a voice synthesizer, a phototypesetter, a digital switching network, and a satellite PDP-11/20 which generates vectors, curves, and characters on a Tektronix 611 storage-tube display.

The greater part of UNIX software is written in the above-mentioned C language⁶. Early versions of the operating system were written in assembly language, but during the summer of 1973, it was rewritten in C. The size of the new system is about one third greater than the old. Since the new system is not only much easier to understand and to modify but also includes many functional improvements, including multiprogramming and the ability to share reentrant code among several user programs, we considered this increase in size quite acceptable.

3. The File system

The most important role of UNIX is to provide a file system. From the point of view of the user, there are three kinds of files: ordinary disk files, directories, and special files.

3.1 Ordinary files

A file contains whatever information the user places on it, for example symbolic or binary (object) programs. No particular structuring is expected by the system. Files of text consist simply of a string of characters, with lines demarcated by the new-line character. Binary programs are sequences of words as they will appear in core memory when the program starts executing. A few user programs manipulate files with more structure; for example, the assembler generates, and the loader expects, an object file in a particular format. However, the structure of files is controlled by the programs which use them, not by the system.

3.2 Directories

Directories provide the mapping between the names of files and the files themselves, and thus induce a structure on the file system as a whole. Each user has a directory of his own files; he may also create subdirectories to contain groups of files conveniently treated together. A directory behaves exactly like an ordinary file except that it cannot be written on by unprivileged programs, so that the system controls the contents of directories. However, anyone with appropriate permission may read a directory just like any other file.

The system maintains several directories for its own use. One of these is the *root* directory. All files in the system can be found by tracing a path through a chain of directories until the desired file is reached. The starting point for such searches is often the root. Another system directory contains all the programs provided for general use; that is, all the *commands*. As will be seen, however, it is by no means necessary that a program reside in this directory for it to be executed.

Files are named by sequences of 14 or fewer characters. When the name of a file is specified to the system, it may be in the form of a *path name*, which is a sequence of directory names separated by slashes "/" and ending in a file name. If the sequence begins with a slash, the search begins in the root directory. The name */alpha/beta/gamma* causes the system to search the root for directory *alpha*, then to search *alpha* for *beta*, finally to find *gamma* in *beta*. *Gamma* may be an ordinary file, a directory, or a special file. As a limiting case, the name "/" refers to the root itself.

A path name not starting with "/" causes the system to begin the search in the user's current directory. Thus, the name *alpha/beta* specifies the file named *beta* in subdirectory *alpha* of the current directory. The simplest kind of name, for example *alpha*, refers to a file which itself is found in the current directory. As another limiting case, the null file name refers to the current directory.

The same non-directory file may appear in several directories under possibly different names. This feature is called *linking*; a directory entry for a file is sometimes called a link. UNIX differs from other systems in which linking is permitted in that all links to a file have equal status. That is, a file does not exist within a particular directory; the directory entry for a file consists merely of its name and a pointer to the information actually describing the file. Thus a file exists independently of any directory entry, although in practice a file is made to disappear along with the last link to it.

Each directory always has at least two entries. The name "." in each directory refers to the directory itself. Thus a program may read the current directory under the name "." without knowing its complete path name. The name ".." by convention refers to the parent of the directory in which it appears, that is, to the directory in which it was created.

The directory structure is constrained to have the form of a rooted tree. Except for the special entries "." and "..", each directory must appear as an entry in exactly one other, which is its parent. The reason for this is to simplify the writing of programs which visit subtrees of the directory structure, and more important, to avoid the separation of portions of the hierarchy. If arbitrary links to directories were permitted, it would be quite difficult to detect when the last connection from the root to a directory was severed.

3.3 Special files

Special files constitute the most unusual feature of the UNIX file system. Each I/O device supported by UNIX is associated with at least one such file. Special files are read and written just like ordinary disk files, but requests to read or write result in activation of the associated device. An entry for each special file resides in directory */dev*, although a link may be made to one of these files just like an ordinary file. Thus, for example, to punch paper tape, one may write on the file */dev/ppt*. Special files exist for each communication line, each disk, each tape drive, and for physical core memory. Of course, the active disks and the core special file are protected from indiscriminate access.

There is a threefold advantage in treating I/O devices this way: file and device I/O are as similar as possible; file and device names have the same syntax and meaning, so that a program expecting a file name as a parameter can be passed a device name; finally, special files are subject to the same protection mechanism as regular files.

3.4 Removable file systems

Although the root of the file system is always stored on the same device, it is not necessary that the entire file system hierarchy reside on this device. There is a *mount* system request which has two arguments: the name of an existing ordinary file, and the name of a special file whose associated storage volume (e. g. disk pack) should have the structure of an independent file system containing its own directory hierarchy. The effect of *mount* is to cause references to the heretofore ordinary file to refer instead to the root directory of the file system on the removable volume. In effect, *mount* replaces a leaf of the hierarchy tree (the ordinary file) by a whole new subtree (the hierarchy stored on the removable volume). After the *mount*, there is virtually no distinction between files on the removable volume and those in the permanent file system. In our installation, for example, the root directory resides on the fixed-head disk, and the large disk drive, which contains user's files, is mounted by the system initialization program; the four smaller disk drives are available to users for mounting their own disk packs. A mountable file system is generated by writing on its corresponding special file. A utility program is available to create an empty file system, or one may simply copy an existing file system.

There is only one exception to the rule of identical treatment of files on different devices: no link may exist between one file system hierarchy and another. This restriction is enforced so as to avoid the elaborate bookkeeping which would otherwise be required to assure removal of the links when the removable volume is finally dismounted. In particular, in the root directories of all file systems, removable or not, the name *..* refers to the directory itself instead of to its parent.

3.5 Protection

Although the access control scheme in UNIX is quite simple, it has some unusual features. Each user of the system is assigned a unique user identification number. When a file is created, it is marked with the user ID of its owner. Also given for new files is a set of seven protection bits. Six of these specify independently read, write, and execute permission for the owner of the file and for all other users.

If the seventh bit is on, the system will temporarily change the user identification of the current user to that of the creator of the file whenever the file is executed as a program. This change in user ID is effective only during the execution of the program which calls for it. The set-user-ID feature provides for privileged programs which may use files inaccessible to other users. For example, a program may keep an accounting file which should neither be read nor changed except by the program itself. If the set-user-identification bit is on for the program, it may access the file although this access might be forbidden to other programs invoked by the given program's user. Since the actual user ID of the invoker of any program is always available, set-user-ID programs may take any measures desired to satisfy themselves as to their invoker's credentials. This mechanism is used to allow users to execute the carefully-written commands which call privileged system entries. For example, there is a system entry invocable only by the "super-user" (below) which creates an empty directory. As indicated above, directories are expected to have entries for *.* and *..*. The command which creates a directory is owned by the super-user and has the set-user-ID bit set. After it checks its invoker's authorization to

create the specified directory, it creates it and makes the entries for "." and ".."

Since anyone may set the set-user-ID bit on one of his own files, this mechanism is generally available without administrative intervention. For example, this protection scheme easily solves the MOO accounting problem posed in [7].

The system recognizes one particular user ID (that of the "super-user") as exempt from the usual constraints on file access; thus (for example) programs may be written to dump and reload the file system without unwanted interference from the protection system.

3.6 I/O calls

The system calls to do I/O are designed to eliminate the differences between the various devices and styles of access. There is no distinction between "random" and "sequential" I/O, nor is any logical record size imposed by the system. The size of an ordinary file is determined by the highest byte written on it; no predetermination of the size of a file is necessary or possible.

To illustrate the essentials of I/O in UNIX, some of the basic calls are summarized below in an anonymous language which will indicate the required parameters without getting into the complexities of machine language programming. Each call to the system may potentially result in an error return, which for simplicity is not represented in the calling sequence.

To read or write a file assumed to exist already, it must be opened by the following call:

```
filep = open ( name, flag )
```

Name indicates the name of the file. An arbitrary path name may be given. The *flag* argument indicates whether the file is to be read, written, or "updated," that is read and written simultaneously.

The returned value *filep* is called a *file descriptor*. It is a small integer used to identify the file in subsequent calls to read, write or otherwise manipulate the file.

To create a new file or completely rewrite an old one, there is a *create* system call which creates the given file if it does not exist, or truncates it to zero length if it does exist. *Create* also opens the new file for writing and, like *open*, returns a file descriptor.

There are no user-visible locks in the file system, nor is there any restriction on the number of users who may have a file open for reading or writing. Although it is possible for the contents of a file to become scrambled when two users write on it simultaneously, in practice difficulties do not arise. We take the view that locks are neither necessary nor sufficient, in our environment, to prevent interference between users of the same file. They are unnecessary because we are not faced with large, single-file data bases maintained by independent processes. They are insufficient because locks in the ordinary sense, whereby one user is prevented from writing on a file which another user is reading, cannot prevent confusion when, for example, both users are editing a file with an editor which makes a copy of the file being edited.

It should be said that the system has sufficient internal interlocks to maintain the logical consistency of the file system when two users engage simultaneously in such inconvenient activities as writing on the same file, creating files in the same directory, or deleting each other's open files.

Except as indicated below, reading and writing are sequential. This means that if a particular byte in the file was the last byte written (or read), the next I/O call implicitly refers to the first following byte. For each open file there is a pointer, maintained by the system, which indicates the next byte to be read or written. If *n* bytes are read or written, the pointer advances by *n* bytes.

Once a file is open, the following calls may be used.

```
n = read ( filep, buffer, count )
```

```
n = write ( filep, buffer, count )
```

Up to *count* bytes are transmitted between the file specified by *filep* and the byte array specified by *buffer*. The returned value *n* is the number of bytes actually transmitted. In the *write* case, *n* is the same as *count* except under exceptional conditions like I/O errors or end of physical medium on special files, in a *read*, however, *n* may without error be less than *count*. If the read pointer is so near the end of the file that reading *count* characters would cause reading beyond the end, only sufficient bytes are transmitted to reach the end of the file; also, typewriter-like devices never return more than one line of input. When a *read* call returns with *n* equal to zero, it indicates the end of the file. For disk files this occurs when the read pointer becomes equal to the current size of the file. It is possible to generate an end-of-file from a typewriter by use of an escape sequence which depends on the device used.

Bytes written on a file affect only those implied by the position of the write pointer and the count; no other part of the file is changed. If the last byte lies beyond the end of the file, the file is grown as needed.

To do random (direct access) I/O it is only necessary to move the read or write pointer to the appropriate location in the file.

location = seek (filep, offset, base)

The pointer associated with *filep* is moved to a position *offset* bytes from the beginning of the file, from the current position of the pointer, or from the end of the file, depending on *base*. *Offset* may be negative. For some devices (e.g. paper tape and typewriters) seek calls are ignored. The actual offset from the beginning of the file to which the pointer was moved is returned in *location*.

3.6.1 Other I/O calls

There are several additional system entries having to do with I/O and with the file system which will not be discussed. For example: close a file, get the status of a file, change the protection mode or the owner of a file, create a directory, make a link to an existing file, delete a file.

4. Implementation of the file system

As mentioned in §3.2 above, a directory entry contains only a name for the associated file and a pointer to the file itself. This pointer is an integer called the *i-number* (for index number) of the file. When the file is accessed, its *i-number* is used as an index into a system table (the *i-list*) stored in a known part of the device on which the directory resides. The entry thereby found (the file's *i-node*) contains the description of the file:

1. its owner,
2. its protection bits;
3. the physical disk or tape addresses for the file contents;
4. its size,
5. time of last modification,
6. the number of links to the file; that is, the number of times it appears in a directory,
7. a bit indicating whether the file is a directory;
8. a bit indicating whether the file is a special file;
9. a bit indicating whether the file is "large" or "small."

The purpose of an *open* or *create* system call is to turn the path name given by the user into an *i-number* by searching the explicitly or implicitly named directories. Once a file is open, its device, *i-number*, and read/write pointer are stored in a system table indexed by the file descriptor returned by the *open* or *create*. Thus the file descriptor supplied during a subsequent call to read or write the file may be easily related to the information necessary to access the file.

When a new file is created, an *i-node* is allocated for it and a directory entry is made which contains the name of the file and the *i-node* number. Making a link to an existing file involves creating a directory entry with the new name, copying the *i-number* from the original file entry, and incrementing the

link-count field of the i-node. Removing (deleting) a file is done by decrementing the link-count of the i-node specified by its directory entry and erasing the directory entry. If the link-count drops to 0, any disk blocks in the file are freed and the i-node is deallocated.

The space on all fixed or removable disks which contain a file system is divided into a number of 512-byte blocks logically addressed from 0 up to a limit which depends on the device. There is space in the i-node of each file for eight device addresses. A *small* (non-special) file fits into eight or fewer blocks; in this case the addresses of the blocks themselves are stored. For *large* (non-special) files, seven of the eight device addresses may point to indirect blocks each containing 256 addresses for the data blocks of the file. If required, the eighth word is the address of a double-indirect block containing 256 more addresses of indirect blocks. Thus files may conceptually grow to $(7+256) \cdot 256 \cdot 512$ bytes, actually they are restricted to $16,777,216 (2^{24})$ bytes. Once opened, a small file (size 1 to 8 blocks) can be accessed directly. A large file (size 9 to 32768 blocks) requires one additional access to read below logical block 1792 ($7 \cdot 256$) and two additional references above 1792.

The foregoing discussion applies to ordinary files. When an I/O request is made to a file whose i-node indicates that it is special, the last seven device address words are immaterial, and the first is interpreted as a pair of bytes which constitute an internal *device name*. These bytes specify respectively a device type and subdevice number. The device type indicates which system routine will deal with I/O on that device; the subdevice number selects, for example, a disk drive attached to a particular controller or one of several similar typewriter interfaces.

In this environment, the implementation of the *mount* system call (§3.4) is quite straightforward. *Mount* maintains a system table whose argument is the i-number and device name of the ordinary file specified during the *mount*, and whose corresponding value is the device name of the indicated special file. This table is searched for each (i-number, device)-pair which turns up while a path name is being scanned during an *open* or *create*: if a match is found, the i-number is replaced by 1 (which is the i-number of the root directory on all file systems), and the device name is replaced by the table value.

To the user, both reading and writing of files appear to be synchronous and unbuffered. That is, immediately after return from a *read* call the data are available, and conversely after a *write* the user's workspace may be reused. In fact the system maintains a rather complicated buffering mechanism which reduces greatly the number of I/O operations required to access a file. Suppose a *write* call is made specifying transmission of a single byte. UNIX will search its buffers to see whether the affected disk block currently resides in core memory; if not, it will be read in from the device. Then the affected byte is replaced in the buffer and an entry is made in a list of blocks to be written. The return from the *write* call may then take place, although the actual I/O may not be completed until a later time. Conversely, if a single byte is read, the system determines whether the secondary storage block in which the byte is located is already in one of the system's buffers; if so, the byte can be returned immediately. If not, the block is read into a buffer and the byte picked out.

The system recognizes when a program has made accesses to sequential blocks of a file, and asynchronously pre-reads the next block. This significantly reduces the running time of most programs while adding little to system overhead.

A program which reads or writes files in units of 512 bytes has an advantage over a program which reads or writes a single byte at a time, but the gain is not immense; it comes mainly from the avoidance of system overhead. A program which is used rarely or which does no great volume of I/O may quite reasonably read and write in units as small as it wishes.

The notion of the i-list is an unusual feature of UNIX. In practice, this method of organizing the file system has proved quite reliable and easy to deal with. To the system itself, one of its strengths is the fact that each file has a short, unambiguous name which is related in a simple way to the protection, addressing, and other information needed to access the file. It also permits a quite simple and rapid algorithm for checking the consistency of a file system, for example verification that the portions of each device containing useful information and those free to be allocated are disjoint and together exhaust the space on the device. This algorithm is independent of the directory hierarchy, since it need only

scan the linearly-organized i-list. At the same time the notion of the i-list induces certain peculiarities not found in other file system organizations. For example, there is the question of who is to be charged for the space a file occupies, since all directory entries for a file have equal status. Charging the owner of a file is unfair in general, since one user may create a file, another may link to it, and the first user may delete the file. The first user is still the owner of the file, but it should be charged to the second user. The simplest reasonably fair algorithm seems to be to spread the charges equally among users who have links to a file. The current version of UNIX avoids the issue by not charging any fees at all.

4.1 Efficiency of the file system

To provide an indication of the overall efficiency of UNIX and of the file system in particular, timings were made of the assembly of a 8848-line program. The assembly was run alone on the machine; the total clock time was 32 seconds, for a rate of 276 lines per second. The time was divided as follows: 66% assembler execution time, 21% system overhead, 13% disk wait time. We will not attempt any interpretation of these figures nor any comparison with other systems, but merely note that we are generally satisfied with the overall performance of the system.

5. Processes and images

An *image* is a computer execution environment. It includes a core image, general register values, status of open files, current directory and the like. An image is the current state of a pseudo-computer.

A *process* is the execution of an image. While the processor is executing on behalf of a process, the image must reside in core, during the execution of other processes it remains in core unless the appearance of an active, higher-priority process forces it to be swapped out to the fixed-head disk.

The user-core part of an image is divided into three logical segments. The program text segment begins at location 0 in the virtual address space. During execution, this segment is write-protected and a single copy of it is shared among all processes executing the same program. At the first 8K byte boundary above the program text segment in the virtual address space begins a non-shared, writable data segment, the size of which may be extended by a system call. Starting at the highest address in the virtual address space is a stack segment, which automatically grows downward as the hardware's stack pointer fluctuates.

5.1 Processes

Except while UNIX is bootstrapping itself into operation, a new process can come into existence only by use of the *fork* system call:

```
processid = fork ( label )
```

When *fork* is executed by a process, it splits into two independently executing processes. The two processes have independent copies of the original core image, and share any open files. The new processes differ only in that one is considered the parent process: in the parent, control returns directly from the *fork*, while in the child, control is passed to location *label*. The *processid* returned by the *fork* call is the identification of the other process.

Because the return points in the parent and child process are not the same, each image existing after a *fork* may determine whether it is the parent or child process.

5.2 Pipes

Processes may communicate with related processes using the same system *read* and *write* calls that are used for file system I/O. The call

```
filep = pipe ( )
```

returns a file descriptor *filep* and creates an inter-process channel called a *pipe*. This channel, like other

open files, is passed from parent to child process in the image by the *fork* call. A *read* using a pipe file descriptor waits until another process writes using the file descriptor for the same pipe. At this point, data are passed between the images of the two processes. Neither process need know that a pipe, rather than an ordinary file, is involved.

Although inter-process communication via pipes is a quite valuable tool (see §6.2), it is not a completely general mechanism, since the pipe must be set up by a common ancestor of the processes involved.

5.3 Execution of programs

Another major system primitive is invoked by

```
execute ( file, arg1, arg2, ... , argn )
```

which requests the system to read in and execute the program named by *file*, passing it string arguments *arg₁*, *arg₂*, ..., *arg_n*. All the code and data in the process using *execute* is replaced from the *file*, but open files, current directory, and inter-process relationships are unaltered. Only if the call fails, for example because *file* could not be found or because its execute-permission bit was not set, does a return take place from the *execute* primitive; it resembles a "jump" machine instruction rather than a subroutine call.

5.4 Process synchronization

Another process control system call

```
processid = wait ( )
```

causes its caller to suspend execution until one of its children has completed execution. Then *wait* returns the *processid* of the terminated process. An error return is taken if the calling process has no descendants. Certain status from the child process is also available.

5.5 Termination

Lastly,

```
exit ( status )
```

terminates a process, destroys its image, closes its open files, and generally obliterates it. When the parent is notified through the *wait* primitive, the indicated *status* is available to the parent. Processes may also terminate as a result of various illegal actions or user-generated signals (§7 below).

6. The Shell

For most users, communication with UNIX is carried on with the aid of a program called the Shell. The Shell is a command line interpreter: it reads lines typed by the user and interprets them as requests to execute other programs. In simplest form, a command line consists of the command name followed by arguments to the command, all separated by spaces:

```
command arg1 arg2 ... argn
```

The Shell splits up the command name and the arguments into separate strings. Then a file with name *command* is sought; *command* may be a path name including the "/" character to specify any file in the system. If *command* is found, it is brought into core and executed. The arguments collected by the Shell are accessible to the command. When the command is finished, the Shell resumes its own execution, and indicates its readiness to accept another command by typing a prompt character.

If file *command* cannot be found, the Shell prefixes the string */bin/* to *command* and attempts again to find the file. Directory */bin* contains all the commands intended to be generally used.

6.1 Standard I/O

The discussion of I/O in §3 above seems to imply that every file used by a program must be opened or created by the program in order to get a file descriptor for the file. Programs executed by the Shell, however, start off with two open files which have file descriptors 0 and 1. As such a program begins execution, file 1 is open for writing, and is best understood as the standard output file. Except under circumstances indicated below, this file is the user's typewriter. Thus programs which wish to write informative or diagnostic information ordinarily use file descriptor 1. Conversely, file 0 starts off open for reading, and programs which wish to read messages typed by the user usually read this file.

The Shell is able to change the standard assignments of these file descriptors from the user's typewriter printer and keyboard. If one of the arguments to a command is prefixed by ">", file descriptor 1 will, for the duration of the command, refer to the file named after the ">". For example,

```
ls
```

ordinarily lists, on the typewriter, the names of the files in the current directory. The command

```
ls >there
```

creates a file called *there* and places the listing there. Thus the argument ">there" means, "place output on *there*". On the other hand,

```
ed
```

ordinarily enters the editor, which takes requests from the user via his typewriter. The command

```
ed <script
```

interprets *script* as a file of editor commands; thus "<script" means, "take input from *script*".

Although the file name following "<" or ">" appears to be an argument to the command, in fact it is interpreted completely by the Shell and is not passed to the command at all. Thus no special coding to handle I/O redirection is needed within each command; the command need merely use the standard file descriptors 0 and 1 where appropriate.

6.2 Filters

An extension of the standard I/O notion is used to direct output from one command to the input of another. A sequence of commands separated by vertical bars causes the Shell to execute all the commands simultaneously and to arrange that the standard output of each command be delivered to the standard input of the next command in the sequence. Thus in the command line

```
ls | pr -2 | opr
```

ls lists the names of the files in the current directory; its output is passed to *pr*, which paginates its input with dated headings. The argument "-2" means double column. Likewise the output from *pr* is input to *opr*. This command spools its input onto a file for off-line printing.

This procedure could have been carried out more clumsily by

```
ls >temp1
pr -2 <temp1 >temp2
opr <temp2
```

followed by removal of the temporary files. In the absence of the ability to redirect output and input, a still clumsier method would have been to require the *ls* command to accept user requests to paginate its output, to print in multi-column format, and to arrange that its output be delivered off-line. Actually it would be surprising, and in fact unwise for efficiency reasons, to expect authors of commands such as *ls* to provide such a wide variety of output options.

A program such as *pr* which copies its standard input to its standard output (with processing) is called a *filter*. Some filters which we have found useful perform character transliteration, sorting of the input, and encryption and decryption.

6.3 Command Separators; Multitasking

Another feature provided by the Shell is relatively straightforward. Commands need not be on different lines; instead they may be separated by semicolons.

```
ls; ed
```

will first list the contents of the current directory, then enter the editor.

A related feature is more interesting. If a command is followed by "&", the Shell will not wait for the command to finish before prompting again; instead, it is ready immediately to accept a new command. For example,

```
as source >output &
```

causes *source* to be assembled, with diagnostic output going to *output*; no matter how long the assembly takes, the Shell returns immediately. When the Shell does not wait for the completion of a command, the identification of the process running that command is printed. This identification may be used to wait for the completion of the command or to terminate it. The "&" may be used several times in a line:

```
as source >output & ls >files &
```

does both the assembly and the listing in the background. In the examples above using "&", an output file other than the typewriter was provided; if this had not been done, the outputs of the various commands would have been intermingled.

The Shell also allows parentheses in the above operations. For example

```
(date; ls) >x &
```

prints the current date and time followed by a list of the current directory onto the file *x*. The Shell also returns immediately for another request.

6.4 The Shell as a Command; Command Files

The Shell is itself a command, and may be called recursively. Suppose file *tryout* contains the lines

```
as source
mv a.out testprog
testprog
```

The *mv* command causes the file *a.out* to be renamed *testprog*. *A.out* is the (binary) output of the assembler, ready to be executed. Thus if the three lines above were typed on the console, *source* would be assembled, the resulting program renamed *testprog*, and *testprog* executed. When the lines are in *tryout*, the command

```
sh <tryout
```

would cause the Shell *sh* to execute the commands sequentially.

The Shell has further capabilities, including the ability to substitute parameters and to construct argument lists from a specified subset of the file names in a directory. It is also possible to execute commands conditionally on character string comparisons or on existence of given files and to perform transfers of control within filed command sequences.

6.5 Implementation of the Shell

The outline of the operation of the Shell can now be understood. Most of the time, the Shell is waiting for the user to type a command. When the new-line character ending the line is typed, the Shell's *read* call returns. The Shell analyzes the command line, putting the arguments in a form appropriate for *execute*. Then *fork* is called. The child process, whose code of course is still that of the Shell, attempts to perform an *execute* with the appropriate arguments. If successful, this will bring in and start execution of the program whose name was given. Meanwhile, the other process resulting from the *fork*, which is the parent process, *waits* for the child process to die. When this happens, the Shell knows the command is finished, so it types its prompt and reads the typewriter to obtain another command.

Given this framework, the implementation of background processes is trivial: whenever a command line contains "&", the Shell merely refrains from waiting for the process which it created to execute the command.

Happily, all of this mechanism meshes very nicely with the notion of standard input and output files. When a process is created by the *fork* primitive, it inherits not only the core image of its parent but also all the files currently open in its parent, including those with file descriptors 0 and 1. The Shell, of course, uses these files to read command lines and to write its prompts and diagnostics, and in the ordinary case its children—the command programs—inherit them automatically. When an argument with "<" or ">" is given however, the offspring process, just before it performs *execute*, makes the standard I/O file descriptor 0 or 1 respectively refer to the named file. This is easy because, by agreement, the smallest unused file descriptor is assigned when a new file is *opened* (or *created*); it is only necessary to close file 0 (or 1) and open the named file. Because the process in which the command program runs simply terminates when it is through, the association between a file specified after "<" or ">" and file descriptor 0 or 1 is ended automatically when the process dies. Therefore the Shell need not know the actual names of the files which are its own standard input and output, since it need never reopen them.

Filters are straightforward extensions of standard I/O redirection with pipes used instead of files.

In ordinary circumstances, the main loop of the Shell never terminates. (The main loop includes that branch of the return from *fork* belonging to the parent process; that is, the branch which does a *wait*, then reads another command line.) The one thing which causes the Shell to terminate is discovering an end-of-file condition on its input file. Thus, when the Shell is executed as a command with a given input file, as in

```
sh <com:file
```

the commands in *comfile* will be executed until the end of *comfile* is reached; then the instance of the Shell invoked by *sh* will terminate. Since this Shell process is the child of another instance of the Shell, the *wait* executed in the latter will return, and another command may be processed.

6.6 Initialization

The instances of the Shell to which users type commands are themselves children of another process. The last step in the initialization of UNIX is the creation of a single process and the invocation (via *execute*) of a program called *init*. The role of *init* is to create one process for each typewriter channel which may be dialed up by a user. The various subinstances of *init* open the appropriate typewriters for input and output. Since when *init* was invoked there were no files open, in each process the typewriter keyboard will receive file descriptor 0 and the printer file descriptor 1. Each process types out a message requesting that the user log in and waits, reading the typewriter, for a reply. At the outset, no one is logged in, so each process simply hangs. Finally someone types his name or other identification. The appropriate instance of *init* wakes up, receives the log-in line, and reads a password file. If the user name is found, and if he is able to supply the correct password, *init* changes to the user's default current directory, sets the process's user ID to that of the person logging in, and performs an *execute* of the Shell. At this point the Shell is ready to receive commands and the logging-in proto-

col is complete.

Meanwhile, the mainstream path of *init* (the parent of all the subinstances of itself which will later become Shells) does a *wait*. If one of the child processes terminates, either because a Shell found an end of file or because a user typed an incorrect name or password, this path of *init* simply recreates the defunct process, which in turn reopens the appropriate input and output files and types another login message. Thus a user may log out simply by typing the end-of-file sequence in place of a command to the Shell.

6.7 Other programs as Shell

The Shell as described above is designed to allow users full access to the facilities of the system, since it will invoke the execution of any program with appropriate protection mode. Sometimes, however, a different interface to the system is desirable, and this feature is easily arranged.

Recall that after a user has successfully logged in by supplying his name and password, *init* ordinarily invokes the Shell to interpret command lines. The user's entry in the password file may contain the name of a program to be invoked after login instead of the Shell. This program is free to interpret the user's messages in any way it wishes.

For example, the password file entries for users of a secretarial editing system specify that the editor *ed* is to be used instead of the Shell. Thus when editing system users log in, they are inside the editor and can begin work immediately; also, they can be prevented from invoking UNIX programs not intended for their use. In practice, it has proved desirable to allow a temporary escape from the editor to execute the formatting program and other utilities.

Several of the games (e.g., chess, blackjack, 3D tic-tac-toe) available on UNIX illustrate a much more severely restricted environment. For each of these an entry exists in the password file specifying that the appropriate game-playing program is to be invoked instead of the Shell. People who log in as a player of one of the games find themselves limited to the game and unable to investigate the presumably more interesting offerings of UNIX as a whole.

7. Traps

The PDP-11 hardware detects a number of program faults, such as references to non-existent memory, unimplemented instructions, and odd addresses used where an even address is required. Such faults cause the processor to trap to a system routine. When an illegal action is caught, unless other arrangements have been made, the system terminates the process and writes the user's image on file *core* in the current directory. A debugger can be used to determine the state of the program at the time of the fault.

Programs which are looping, which produce unwanted output, or about which the user has second thoughts may be halted by the use of the *interrupt* signal, which is generated by typing the "delete" character. Unless special action has been taken, this signal simply causes the program to cease execution without producing a core image file.

There is also a *quit* signal which is used to force a core image to be produced. Thus programs which loop unexpectedly may be halted and the core image examined without prearrangement.

The hardware-generated faults and the interrupt and quit signals can, by request, be either ignored or caught by the process. For example, the Shell ignores quits to prevent a quit from logging the user out. The editor catches interrupts and returns to its command level. This is useful for stopping long printouts without losing work in progress (the editor manipulates a copy of the file it is editing). In systems without floating point hardware, unimplemented instructions are caught and floating point instructions are interpreted.

8. Perspective

Perhaps paradoxically, the success of UNIX is largely due to the fact that it was not designed to meet any predefined objectives. The first version was written when one of us (Thompson), dissatisfied with the available computer facilities, discovered a little-used PDP-7 and set out to create a more hospitable environment. This essentially personal effort was sufficiently successful to gain the interest of the remaining author and others, and later to justify the acquisition of the PDP-11/20, specifically to support a text editing and formatting system. When in turn the 11/20 was outgrown, UNIX had proved useful enough to persuade management to invest in the PDP-11/45. Our goals throughout the effort, when articulated at all, have always concerned themselves with building a comfortable relationship with the machine and with exploring ideas and inventions in operating systems. We have not been faced with the need to satisfy someone else's requirements, and for this freedom we are grateful.

Three considerations which influenced the design of UNIX are visible in retrospect.

First: since we are programmers, we naturally designed the system to make it easy to write, test, and run programs. The most important expression of our desire for programming convenience was that the system was arranged for interactive use, even though the original version only supported one user. We believe that a properly-designed interactive system is much more productive and satisfying to use than a "batch" system. Moreover such a system is rather easily adaptable to non-interactive use, while the converse is not true.

Second: there have always been fairly severe size constraints on the system and its software. Given the partially antagonistic desires for reasonable efficiency and expressive power, the size constraint has encouraged not only economy but a certain elegance of design. This may be a thinly disguised version of the "salvation through suffering" philosophy, but in our case it worked.

Third: nearly from the start, the system was able to, and did, maintain itself. This fact is more important than it might seem. If designers of a system are forced to use that system they quickly become aware of its functional and superficial deficiencies and are strongly motivated to correct them before it is too late. Since all source programs were always available and easily modified on-line, we were willing to revise and rewrite the system and its software when new ideas were invented, discovered, or suggested by others.

The aspects of UNIX discussed in this paper exhibit clearly at least the first two of these design considerations. The interface to the file system, for example, is extremely convenient from a programming standpoint. The lowest possible interface level is designed to eliminate distinctions between the various devices and files and between direct and sequential access. No large "access method" routines are required to insulate the programmer from the system calls; in fact all user programs either call the system directly or use a small library program, only tens of instructions long, which buffers a number of characters and reads or writes them all at once.

Another important aspect of programming convenience is that there are no "control blocks" with a complicated structure partially maintained by and depended on by the file system or other system calls. Generally speaking, the contents of a program's address space are the property of the program, and we have tried to avoid placing restrictions on the data structures within that address space.

Given the requirement that all programs should be usable with any file or device as input or output, it is also desirable from a space-efficiency standpoint to push device-dependent considerations into the operating system itself. The only alternatives seem to be to load routines for dealing with each device with all programs, which is expensive in space, or to depend on some means of dynamically linking to the routine appropriate to each device when it is actually needed, which is expensive either in overhead or in hardware.

Likewise, the process control scheme and command interface have proved both convenient and efficient. Since the Shell operates as an ordinary, swappable user program, it consumes no wired-down space in the system proper, and it may be made as powerful as desired at little cost. In particular, given the framework in which the Shell executes as a process which spawns other processes to perform

commands, the notions of I/O redirection, background processes, command files, and user-selectable system interfaces all become essentially trivial to implement.

8.1 Influences

The success of UNIX lies not so much in new inventions but rather in the full exploitation of a carefully selected set of fertile ideas, and especially in showing that they can be keys to the implementation of a small yet powerful operating system.

The *fork* operation, essentially as we implemented it, was present in the Berkeley time sharing system⁸. On a number of points we were influenced by Multics, which suggested the particular form of the I/O system calls⁹ and both the name of the Shell and its general functions. The notion that the Shell should create a process for each command was also suggested to us by the early design of Multics, although in that system it was later dropped for efficiency reasons. A similar scheme is used by TENEX¹⁰.

9. Statistics

The following numbers are presented to suggest the scale of our operation. Those of our users not involved in document preparation tend to use the system for program development, especially language work. There are few important "applications" programs.

Overall, we have

100	user population
14	maximum simultaneous users
380	directories
4800	files
66300	512-byte secondary storage blocks used

There is a "background" process that runs at the lowest possible priority; it is used to soak up any idle CPU time. It has been used to produce a million-digit approximation to the constant $e-2$, and is now solving all rook-and-pawn vs. rook chess endgames. Not counting this background work, we average daily


2400	commands
5.5	CPU hours
100	connect hours
32	different users
100	logins


Acknowledgements. We are grateful to R.H. Canaday, L.L. Cherry, and L.E. McMahon for their contributions to UNIX. We are particularly appreciative of the inventiveness, thoughtful criticism, and constant support of R. Morris, M.D. McIlroy, and J.F. Ossanna.

References

1. Digital Equipment Corporation. *PDP-11/40 Processor Handbook* (1972), *PDP-11/45 Processor Handbook* (1971), and *PDP-11/70 Processor Handbook* (1975).
2. Deutsch, L.P., and Lampson, B.W. An online editor. *Comm. ACM* 10, 12 (Dec. 1967), 793-799, 803.
3. Richards, M. BCPL: A tool for compiler writing and system programming. Proc. AFIPS 1969 SJCC, Vol. 34, AFIPS Press, Montvale, N.J., pp. 557-566.
4. McClure, R.M. TMG—A syntax directed compiler. Proc. ACM 20th Nat. Conf., ACM, 1965, New York, pp. 262-274.
5. Hall, A.D. The M6 macroprocessor. Computing Science Tech. Rep. #2, Bell Telephone Laboratories, 1969.
6. Ritchie, D.M. C reference manual. Unpublished memorandum, Bell Telephone Laboratories

- (1973).
7. Aleph-null. Computer Recreations. *Software Practice and Experience* 1, 2 (Apr.-June 1971), 201-204.
 8. Deutch, L.P. and Lampson, B.W. SDS 930 time-sharing system preliminary reference manual. Doc. 30.10.10, Project GENIE, Univ. Cal. at Berkeley (Apr. 1965).
 9. Feiertag, R.J., and Organick, E.I. The Multics input-output system. Proc. Third Symposium on Operating Systems Principles. Oct. 18-20, 1971, ACM, New York, pp. 35-41.
 10. Bobrow, D.G., Burchfiel, J.D., Murphy, D.L., and Tomlinson, R.S. TENEX, a paged time sharing system for the PDP-10. *Comm. ACM* 15., 3 (March 1972) 135-143.





The Unix Time-sharing System -- A Retrospective*

Dennis M. Ritchie

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

Unix is a general-purpose, interactive time-sharing operating system for the DEC PDP-11 series of computers. Since its development about five years ago, it has become quite widely used. This paper will discuss the strong and weak points of Unix and some areas where we have made no effort. The following points are touched on.

- The structure of files: a uniform, randomly addressable sequence of bytes. The irrelevance of the notion of 'record.' The efficiency of the addressing of files.
- The structure of file system devices; directories and files.
- I/O devices are integrated into the file system.
- The user interface: fundamentals of the shell, I/O redirection, and pipes.
- The environment of processes: system calls, signals, and the address space.
- Reliability: crashes, losses of files.
- Security: protecting data from corruption and inspection; protecting the system from stoppages.
- Use of a high-level language - the benefits and the costs.
- What Unix does not do: 'real-time', interprocess communication, asynchronous I/O.
- Recommendations to system designers

Unix is a general-purpose, interactive time-sharing operating system for the DEC PDP-11 series of computers. Since its development about five years ago, it has become quite widely used, although publicity efforts in its behalf have been minimal, and the license under which it is made available explicitly excludes maintenance. Currently it is used by more than 150 universities, fifteen commercial and government organizations (most running several systems), and at about 250 Bell System installations. It is useful on a rather broad

range of configurations, ranging from a large PDP-11/70 supporting 50 users, to a single-user LSI-11 system.

Some general observations

In most ways Unix is a very conservative system. Only a handful of its ideas are genuinely new. In fact, a good case can be made that it is in essence a modern implementation of MIT's CTSS system [Crisman]. This claim is intended as a compliment to both Unix and CTSS. Today, more than fifteen years after CTSS was born, few of the interactive systems we know of are superior to it in ease of use; many are inferior in basic design.

*This paper was presented at the Tenth Hawaii International Conference on the System Sciences, Honolulu, January, 1977.

Unix was never a 'project;' it was not designed to meet any specific need except that felt by its major author, Ken Thompson, and latterly by the author of this paper, for a pleasant environment in which to write and use programs. Although it is rather difficult, after the fact, to try to account for its success, the following reasons seem most important.

- It is simple enough to be comprehended, yet powerful enough to do most of the things its users want.
- The user interface is clean and relatively surprise-free. It is also terse to the point of being cryptic until understood.
- It runs on a machine which has become very popular in its own right.
- Besides the operating system and its basic utilities, a good deal of interesting software is available, including a sophisticated text-processing system which handles complicated mathematical material [Kernighan and Cherry], producing output on a typesetter or a typewriter terminal, and a LALR parser-generator [Aho and Johnson].

This paper discusses the strong and weak points of the system and lists some areas where no effort has been made. Only enough design details are given to motivate the discussion; for more see [Ritchie and Thompson].

The structure of files

The Unix file system is simple in structure; nevertheless it is more powerful and general than those often found even in considerably larger operating systems. Every file is regarded as a featureless, randomly addressable sequence of bytes. The system conceals physical properties of the device on which the file is stored, such as the size of a disk track. The size of a file is the number of bytes it contains; the last byte is determined by the high-water mark of writes to the file. It is not necessary, or in fact possible, to preallocate space for a file. The system calls to read and write each come in only one form, which specifies the local name of an open file, a buffer to or from which to perform I/O, and a byte count. I/O is normally sequential, so the

first byte referred to by a read or write operation immediately follows the final byte affected by the preceding operation. 'Random access' is accomplished using a 'seek' system call, which moves the system's internal offset for the instance of the open file to another location which the next read or write will implicitly address. All I/O appears completely synchronous; read-ahead and write-behind are performed by the system.

This particularly simple way of viewing files was suggested by the Multics I/O system [Feiertag and Organick].

The addressing mechanism for files must be carefully designed if it is to be efficient. Files can be large (16 million bytes), are grown without preallocation, and are randomly accessible. The overhead per file must be small, since there can be many files (we have 7922 on the disk storing most user's files); many of them are small (77 per cent have eight or fewer 512-byte blocks, and 37 per cent are only one block long). The details of the mechanism are given in [Ritchie and Thompson]. Briefly, for each file up to eight disk addresses are stored. If the file is small enough, they are the physical addresses of up to eight 512-byte blocks; if the file is large, they are the addresses of up to seven 'indirect blocks' each of which contains the addresses of up to 256 data blocks. If the file is 'huge,' the eighth basic disk address points to a double-indirect block with 256 addresses of indirect blocks; each of these in turn contains the addresses of 256 data blocks. The eight basic addresses are always kept in memory while the file is open. Thus from one to three disk accesses per block must be made to do I/O on a file, depending on the size of the file. The system, however, maintains a cache of recently accessed blocks; when a file of any size is read or written sequentially the expected number of disk accesses per data block is nearly unity.

No careful study has been made of the efficiency of disk I/O, but a simple experiment indicates that the efficiency is comparable to two other systems, DEC's IAS for the PDP-11, and Honeywell's GCOS TSS system running on the H6070. The experiment consisted of timing a program which copied a file which, on the PDP-11's, contained 480 blocks (245,760 bytes). The file

Input and output are ordinarily synchronous; programs wait until their I/O is completed. For disk files, read-ahead and write-behind are handled by the operating system. The mechanisms are efficient enough, and the simplification in user-level code large enough, that we have no general doubts about the wisdom of doing things in this way. There remain special applications in which one desires to initiate I/O on several streams and delay until the operation is complete on only one of them. When the number of streams is small it is possible to simulate this usage with several processes. However, the writers of a Unix NCP ('network control program') interface to the ARPANET [Chesson] feel that genuinely asynchronous I/O would improve their implementation significantly.

Memory is not shared between processes, except for the (read-only) program text. Partly to alleviate the sometimes severe restrictions on the virtual address space imposed by the PDP-11, and partly to simplify communication among tightly-coupled but asynchronous processes, the controlled sharing of writable data areas would be valuable to some applications. The relatively limited virtual address space available on the PDP-11 turns out to be of particular importance. A number of projects which use Unix as a base desire better inter-process communication (both by means of messages and by sharing memory) because they are driven to use several processes for a task which logically requires only one. This is true of several Bell System applications; as its description makes clear, it is also true of INGRES [Stonebraker et al.].

Unix does not attempt to assign non-sharable devices to users. Some devices can only be opened by one process, but there is no mechanism for reserving devices for a particular user for a period of time or over several commands. Few installations with which we have communicated feel this to be a problem. The line printer, for example, is usually dedicated to a spooling program, and its direct use is either forbidden or managed informally. Tapes are always allocated informally. Should the need arise, however, it is worth noting that commands to assign and release devices may be implemented without changing the operating system. Since the

same protection mechanism applies to device files as to ordinary files, an 'assign' command would operate essentially by changing the owner identification attached to the requested device to that of the invoker.

Recommendations

The following points are earnestly recommended to designers of operating systems.

There is really no excuse for not providing a hierarchically-arranged file system. It is very useful to maintain directories containing related files, it is efficient because the amount of searching for files is bounded, and it is easy to provide.

The notion of 'record' seems to be an obsolete remnant of the days of the 80-column card. Files should consist of a sequence of bytes.

The greatest care should be taken to assure that there is only one format for files. This is essential for making programs work smoothly together.

Systems should be written in a high-level language which encourages portability. Manufacturers who build more than one line of machines and also build more than one operating system and set of utilities are wasting money.

Acknowledgement

Much, even most, of the design and implementation of Unix is the work of Ken Thompson. My use of the term 'we' in this paper is intended to include him; I hope his views have not been misrepresented.

References

Aho, A. V., and Johnson, S. C. LR Parsing. *Computing Surveys* 6, 2 (June 1974). pp. 99-124.

Bayer, D. L., and Lycklama, H. MERT - A Multi-environment Real-time Operating System. *Proc. Fifth Symposium on Operating System Principles*. Nov. 19-21, 1975. ACM, New York, pp. 33-42.

Crisman, P. A. (Ed.). *The Compatible Time-sharing System*. MIT Press, 1965.

Chesson, G. L. *The Network Unix System*. *Proc. Fifth Symposium on Operating System Principles*. Nov. 19-21, 1975.

ACM, New York, pp. 60-66.

Dolotta, T. A., and Mashey, J. R. An Introduction to the Programmer's Workbench. Proc. Second International Conference on Software Engineering. Oct. 13-15, 1976. IEEE, New York. pp. 164-168.

Kernighan, B. W., and Cherry, L. L. A System for Typesetting Mathematics. Comm. ACM 18, 3 (Mar. 1975), pp. 151-156.

Kernighan, B. W., and Mashey, J. R. Software Reliability. To appear.

Feiertag, R. J., and Organick, E. I. The Multics Input-output System. Proc. Third Symposium on Operating Systems Principles. Oct. 18-20, 1971. ACM, New York, pp. 18-20.

Mashey, J. R. Using a Command Language as a High-level Programming Language. Proc. Second International Conference on Software Engineering. Oct. 13-15, 1976. IEEE, New York, pp. 169-176.

Ritchie, D. M., Kernighan, B. W., and Lesk, M. E. The C Programming Language. Comp. Sci. Tech. Rep. #31, Bell Laboratories, 1975.

Ritchie, D. M. and Thompson, K. The UNIX Time-sharing System. Comm. ACM 17, 7 (Jul. 1974), pp. 365-375.

Stonebraker, M., Wong, E., Kreps, P., and Held, G. The Design and Implementation of INGRES. ACM Trans. on Database Systems 1, 3 (Sep. 1976). pp. 189-222.

and between actual and formal function arguments) are detected. In some cases this is due to laziness on the part of the compiler, but in others the ability to pun is used explicitly. For example, it is possible to write a general storage allocator which accepts a number of bytes and returns a pointer to an area of the requested size; then an object of appropriate type may be stored in the space referred to by the pointer.

Some way of escaping from the rigor of a strongly typed language is necessary at times in a system programming language, but C is clearly too permissive. The problem lies not only in the possibility of programming errors which could potentially be detected at compile time; almost as important, the portability of programs written in the language suffers. For example, in current C implementations, it is possible to assign a pointer to an integer and then back to a pointer without diagnostics and in fact without changing the value of the resulting pointer. But for this to be possible implies that integers and pointers are the same size, which is an unwarranted assumption in general.

To gain the benefits of a strongly typed language without losing the ability to cheat occasionally and escape the type rules, we have implemented a verifier program which reads a collection of C programs and reports not only on syntax errors but also on violation of type rules which are not checked by the compiler itself. Since the verifier will treat many program files simultaneously, it also checks consistency of usage of separately-compiled procedures, which the compiler itself would be hard-pressed to do.

The Unix system and its central software were originally written in assembly language, before C was invented, and have since been rewritten in C and extended considerably. Thus we are in a better position than most to gauge the effect of using a higher-level language on writing systems. Briefly, the effects were remarkably beneficial and the costs minuscule by comparison. The effects cannot be quantized, since we do not measure productivity by lines of code, but it is suggestive to say that Unix offers a good deal of interesting software, ranging from parser-generators

through mathematical equation formatting packages, which would never have been written at all if their authors had had to write assembly code; many of our most inventive contributors do not know, and do not wish to learn, the instruction set of the PDP-11.

The C versions of programs which existed before C became available are much more easily understood, repaired, and extended than the old versions. This applies especially to the operating system itself. The original system was very difficult to modify, especially to add new devices, but even to make minor changes. The C version is readily modifiable by comparison, and not only by us; more than one university, for example, has completely rewritten the typewriter device driver to suit their own design.

One of the most exciting possibilities for a system written in a higher-level language is transportation of the system to another machine. C programs tend to be portable, despite the problems mentioned above, and we are already reaping benefits from this fact. C compilers (written on Unix) are available at the major Bell Labs computation centers, which have equipment both from Honeywell (running GCOS), and from IBM (running OS, TSO, and TSS). Many major programs written on Unix, including the C compilers themselves, the typesetting software, the equation formatter, the parser-generator, and the text editor, have been moved to one or both of the Honeywell and IBM machines. It has become clear that the most difficult problem in moving this software is not in expressing the algorithms in portable C, but in the system interface. Thus, if the operating system itself were portable, the portability of software would at a stroke be improved considerably. It appears that the job is not so difficult as it might seem. The bulk of the system consists, in fact, of machine-independent code expressing policies and algorithms. Even the file system is quite machine-independent, given a few primitives for communicating with user programs and with devices.

The transportation of Unix itself has not been accomplished, but is the subject of our current research. It successful comple-

tion will free us from dependence on a particular manufacturer or machine architecture. We hope to produce an operating system and set of software which runs on several machines and whose expression in source code is, except for a few modules, identical on each machine. The payoff from such a system, either to an organization which uses several kinds of hardware or to a manufacturer who produces more than one line of machines, should be evident.

Compared to the benefits, the costs of using a higher-level language seem negligible. Certainly the object programs written by the compiler are somewhat larger than those that would be produced by a careful assembly-language coder. It is hard to estimate the average increase in size, because in rewriting it is hard to resist the chance to redesign somewhat (and usually improve). A typical inflation factor would be about 20%. The decrease in speed is likely to be a bit larger, mainly because subroutine linkage tends to be more costly in C (and other high-level languages) than in assembler. However, it is by now a matter of common knowledge that a tiny fraction of the code of any program is likely to consume most of the time, and our experience certainly confirms this belief. A profiling tool for C programs has been useful in making those programs that need to be efficient acceptably fast.

The cost of the compilations themselves has to be considered. This too we deem acceptable. For example, to compile and link-edit the entire operating system ('sysgen') takes a bit over nine minutes of clock time (of which seven minutes are CPU time); the system consists of about 12,000 lines of C code, leading to a rate of about 22 lines per second from source to executable object on a PDP-11/45 with Fabritek cache memory. The compiler is faster than this figure would indicate; the system source makes heavy use of 'include' files, so the number of lines processed by the compiler is actually 31.718 and the rate is 57 lines per second.

These days all the best authorities advocate the use of a high-level language, so Unix can hardly be accused of starting a revolution with this as goal. Still, not all of those who actually produce systems have

leaped on the bandwagon. Perhaps Unix can provide the required nudge. In its largest configurations, it serves 50 simultaneous users (which is about twice the number the manufacturer's most nearly comparable system handles); in a somewhat cut-down version, still written in C and still recognizably the same system, it occupies 8K words and supports a single user on the LSI-11 micro-computer.

What Unix doesn't do

A number of facilities provided in various other systems are not present in Unix. Many of these things would be useful, or even vital, to some applications, and we do not necessarily view their lack as one of the strengths of the system. A good case can be made for the inclusion of several of the following ideas in standard Unix, if only because they have in fact been implemented by others, often in incompatible ways.

Unix is not a 'real-time' system in the sense that it is not possible to lock a process in memory so as to guarantee rapid response to events, nor to connect directly to I/O devices. MERT [Bayer and Lycklama], which is in a sense a generalization of Unix, does allow these operations and in fact all those mentioned in this section. It is a three-level system, with a kernel, several supervisor processes, and user processes. One of the standard supervisor processes is a 'Unix emulator, so that all the ordinary Unix software is available.

There is no general inter-process message facility, nor even a limited communication scheme such as semaphores. It turns out that the pipe mechanism mentioned above is sufficient to implement whatever communication is needed between closely-related, cooperating processes; 'closely-related' means processes with a common ancestor which sets up the channels. Pipes are not, however, of any use in communicating with daemon processes set up to serve several users. At one of the sites at which Unix is run a scheme of 'named pipes' has been implemented. This involves a named file read by a single process which delays until information messages are written into the file by anyone (with permission) who cares to send a message.

The discussion under the heading of Security points out that overconsumption of resources other than swap space occurs, but generally does not cause a crash, although the system may not be very useful for a period. In most such cases a really general remedy is hard to imagine. For example, if one insists on using almost all of the file storage space for storing files, one is certain to run out now and then, and a quota system is unlikely to be of much help, since the space is almost certainly overallocated. An automatically enforced file-space quota would help, however, in the case of the user who accidentally creates a monstrous file.

Hardware is by far the most frequent cause of crashes, and in a basically healthy machine, the most frequent difficulty is momentary power dips, which tend to cause disks to go off line and the processor to enter peculiar, undocumented states. Other kinds of failures occur less often. It does seem characteristic of the PDP-11, particularly in large configurations, to develop transient, hard-to-diagnose Unibus maladies.

A reliable system should not lose or corrupt user's files. The operating system does not take any unusual precautions in this regard. Data destined to be written on the disk may remain in an associative memory cache for up to 15 seconds. Nevertheless our own machine has ruined only three or four files in the past year, not counting files being created at the time of a crash. The rate of destruction of files by the system is negligible compared to that by users accidentally removing or overwriting them. Dumps at near-daily intervals guard against disaster (which has occurred—in historical memory we have had head crashes, and twice a sick disk controller began writing garbage instead of what was asked.)

The reliability of user-level software provided with the system is also good. Many of the reasons for this are related to the positive aspects of the system discussed elsewhere in this paper. Almost all software is written in a high-level language, it tends to be relatively simple and modular with clean interfaces to the operating system and to the outside world, and it tends to be small. For example, the only programs which exceed 5,000 lines of code are the resident operating system itself, the C com-

piler, and the phototypesetter formatter. Of the system 'utilities' (editor, command interpreter, debugger, file maintenance, etc., but not compilers) the upper quartile in number of source lines is 409; the median is 205. The total number of lines in all these programs is about 50,000.

Kernighan and Mashey discuss Unix software reliability more fully than can be done here.

Security

'Security' means the ability to protect against unwanted accessing or destruction of data and against denial of service to others, for example by causing a crash. Because we have always wanted to run a fairly open system, the continuous, close attention required to maintain a secure system has not been paid; as a result there are several security problems.

The weakest area is in protecting against crashing, or at least crippling, the operation of the system. The standard version lacks checks for overconsumption of certain resources, like file space, total number of files, and number of processes (they are limited on a per-user basis in variant versions). Running out of these things does not cause a crash, but will make the system unusable for a period. When resource exhaustion occurs, it is generally evident what happened and who was responsible, so malicious actions are detectable, but the real problem is the accidental program bug.

The theoretical aspects of the situation are brighter in the area of information protection. Each file is marked with its owner and the 'group' of users to which the owner belongs. Files also have a set of nine protection bits divided into three sets of three bits specifying permission to read, to write, or execute as a program. The three sets indicate the permissions applicable to the owner of the file, to members of the owner's group, and to all others.

For directories, the meaning of the access bits are modified: 'read' means the ability to read as a file, that is find out all names in the directory; 'execute' means the ability to search a directory for a given name when it appears as part of a qualified name;

'write' means ability to create and delete files in that directory, and is unrelated to writing of files named.

This classification is not fine enough to account for the needs of all users, but is in general adequate. In fact, most installations do not use groups at all (all users are in the same group), and even those that do would be happy to have more possible user-IDs and fewer group IDs (there are currently 256 of each).

One particular user (the 'super-user') is able to access all files without regard to permissions. This user is also the only one permitted to exercise privileged system entries.

An unusual feature of the protection system is the 'set-user-ID' bit. When this protection bit is on for a file, and the file is executed as a program, the user number used in file permission checking is not that of the person running the program but that of the owner of the file. In practice the bit is used to mark the programs which perform the privileged system functions mentioned above (such as creation of directories, changing the owner of a file, and so forth).

In theory the protection scheme is adequate to maintain security, but in practice breakdowns can easily occur. Most often these come from incorrect protection modes on files. Our software tends to create files accessible, even writable, by everyone. This is not an accident, but a reflection of the open environment in which we operate. Nevertheless, people in more hostile situations must adjust modes frequently; it is easy to forget, and in any case there are brief periods when the modes are wrong. It would be better if software created files in a default mode specifiable by each user. The system administrators must be even more careful than the users to apply proper protection. For example, it is easy to write a program which interprets the contents of a special file referring to a disk and reads or writes the files stored on it, if it is not protected. If a set-user-ID file is writable, another user can copy his own program on top of it.

It is also possible to take advantage of bugs in privileged set-user-ID programs. For example, the program which sends mail

to other users might be given the ability to send to directories which are otherwise protected. It must be carefully written in order to avoid being fooled into mailing other people's private files to its invoker.

There are thus a number of practical difficulties in maintaining a fully secure system. However, the operating system itself seems capable of maintaining data security. The word 'seems' must be used because the system has not been formally verified, yet no security-relevant bugs are known (except the ability to run it out of resources, which was mentioned above). In some ways, in fact, Unix is inherently safer than many other systems. For example, I/O is always done on open files, which are named by an object local to a process. Permissions are checked when the file is opened. The I/O calls themselves have as argument only the (local) name of the open file, and the specification of the user's buffer; physical I/O occurs to a system buffer, and the data is copied in or out of the user's address space by a single place in the system. Thus there is no need for a complicated, bug-prone verification of device commands and channel programs supplied by the user. Likewise, the absence of user 'data control blocks' or other control blocks from the user's address space means that the interface between user processes and the system is rather easily checked, since is conducted by means of explicit arguments.

Use of a Higher-level Language

Essentially all of Unix and its software are written in the C language [Ritchie], which was developed under Unix. C incorporates most of the modern facilities for imposing good structure on control flow and on data and has a rich set of operators. It is a typed language; its fundamental types consist of characters (bytes), integers, long integers, unsigned integers, and single and double precision floating-point quantities. Derived types consist of functions which deliver objects of a given type, arrays of objects of the same type, structures (records) of objects of various types, and pointers to objects of a given type. Standard C is not, however, a strongly-typed language, in that not all type mismatches (as might occur, for example, in assignments

an open file connecting two processes; information written in one end of the pipe may be read from the other, with synchronization, scheduling, and buffering handled by the system. A linear array of processes (a 'pipeline') thus becomes a set of coroutines simultaneously processing an I/O stream. The shell notation for a pipeline separates the names of the various programs by a vertical bar, so for example

```
anycommand | sort | print
```

takes the output of *anycommand*, sorts it, and prints the result in paginated form. The ability to interconnect programs in this way has substantially changed our way of writing and thinking of utility programs in general, and especially those involved with text processing. As a dramatic example, we had three existing programs which would respectively translate characters, sort a file casting out duplicate lines, and compare two sorted files and publish lines which were in one but not the other. Combining these with our on-line dictionary gave a pipeline which would print all words in a document not appearing in the dictionary; in other words, potential spelling mistakes. A simple program to generate plausible derivatives of dictionary words completed the job.

The shell is discussed further by [Mashey] and by [Kernighan and Mashey].

Other aspects of the system, less closely tied to a particular program, are also important. One thing which seems trivial, yet makes a surprising difference once one is used to it, is full-duplex typewriter I/O together with read-ahead. Even though programs generally communicate with the user in terms of lines, not single characters, full-duplex typewriter I/O means that the user can type at any time, even if the system is typing back, without fear of losing or garbling characters. With read-ahead, one need not wait for a response to every line. A good typist entering a document becomes incredibly frustrated at having to pause before starting each new line; for anyone who knows what he wants to say any slowness in response becomes psychologically magnified if the information must be entered bit-by-bit instead of directly.

Both input and output of Unix programs tends to be very terse. This can be

disconcerting to the beginner especially. The editor, for example, has essentially only one diagnostic, namely '?', which means 'you have done something wrong.' Once one knows the editor, the error or difficulty is usually obvious, and the terseness is appreciated after a period of acclimation, but certainly people can be confused at first. However, even if some fuller diagnostics might on occasion be appreciated, there is much noise produced in other systems that we are happy to be rid of. The command interpreter does not remark loudly that each program finished normally, or announce how much space or time it took; the former fact is whispered by an unobtrusive prompt, and anyone who wishes to know the latter may ask explicitly.

Likewise, commands seldom prompt for missing arguments but instead, if the argument is not optional, give at most a one-line summary of their usage and terminate. We know of some systems which seem so proud of their ability to interact that they force interaction on the user whether it is wanted or not. Prompting for missing arguments is an issue of taste which can be discussed in calm tones; insistence on asking questions may cause raised voices.

Although the terseness of typical Unix programs is to some extent a matter of taste, it is also connected with the way programs tend to be combined. A simple example should make the situation clear. The command 'who' writes out one line for each user logged into the system, giving a name, a typewriter name, and the time of login. The command 'wc' (for 'word count') writes out the number of lines, the number of words, and the number of characters in its input. Thus

```
who | wc
```

tells in the line-count field how many users there are. If 'who' produced extraneous verbiage, the count would be off. Worse, if 'wc' insisted on reading from its input which of lines, words, or characters was wanted, it could not be used in this pipeline. Now clearly every command which generates a table should not omit headings; nevertheless we have good reasons to interpret the phrase 'extraneous verbiage' rather liberally.

The process environment

The virtual address space of a process is divided into three regions: a read-only, shared program text region, a writable data area which may grow at one end by explicit request, and a stack which grows at its end automatically as information is pushed on it. The address space contains no 'control blocks' used by the system.

New processes are created by a 'fork' operation, which creates a child process whose code and data are copied from the parent. The child inherits the open files of the parent, and executes asynchronously. The fork mechanism is essential to the basic operation of the system, since each command executed by the shell runs in its own process. This scheme makes a number of services extremely easy to provide. I/O redirection, in particular, is a basically simple operation; it is performed entirely in the subprocess which executes the command, and thus no memory in the parent command interpreter is required to rescind the change in standard input and output. Asynchronous (background) processes likewise require no new mechanism; the shell merely refrains from waiting for the completion of a command specified to be asynchronous. Finally, recursive use of the shell to interpret filed sequences of commands is in no way a special operation.

Communication by processes with the outside world is restricted to a few paths. Explicit system calls, mostly to do I/O, are the most common. A new process receives a set of character-string arguments from its parent upon creation, and returns a byte of status information when it terminates. It may be sent of several 'signals,' which ordinarily force termination, but may, at the choice of the process, be ignored or cause a simulated hardware interrupt. Interrupts from the typewriter, for example, cause a signal to be sent to the processes attached to that typewriter; faults such as addressing errors are also turned into signals. Unassigned signals may be used for communication between cooperating processes. A final, rather specialized, mechanism allows a parent process to trace the actions of its child, receiving notification of faults incurred and accessing the memory of the child. This is used for debugging

There is thus no general inter-process communication or synchronization scheme. As discussed below, this is a weakness of the system, but it is not felt to be important in the uses to which Unix is generally put. Semaphores, for example, can be implemented by using creation and deletion of a known file to represent the P and V operations. Using a semaphore would certainly be more efficient if the notion were made a primitive, but here, as in other aspects of the design, we have preferred to avoid putting new mechanisms in the system which can already be implemented using existing mechanisms. Only when serious and demonstrable inefficiency results is it worth complicating the basic interfaces.

Reliability

The reliability of a system is measured by the absence of unplanned outages, its ability to retain filed information, and the correct functioning of its software.

First, the operating system should not crash. Unix in its present state is commendable in this regard. The six 'Programmer's Workbench' Unix systems [Dolotta] at Bell Laboratories have not had a system failure (outside of announced test periods) attributable to software since November, 1975. Our own system, which changes much more often, has a slightly less perfect record, but crashes are still unusual; it has run for five weeks without being rebooted for any reason. In the past year, about five crashes occurred because of traced, and subsequently repaired bugs; about five more were examined but could not be definitely attributed either to hardware or software.

Two events—running out of swap space, and an unrecoverable I/O error during swapping—cause the system to crash 'voluntarily,' that is, not as a result of a bug causing a fault. It turns out to be rather inconvenient to arrange a more graceful exit for a process which cannot be swapped. Occurrence of swap-space exhaustion can be made arbitrarily rare by providing enough space, and I/O errors in swapping are usually a signal that the hardware is badly impaired, so in neither of these cases do we feel strongly motivated to alleviate the theoretical problems.

stored on another device. It is common for the totality of user files to be too voluminous for a given device. It is then impossible for the directories of all users to be members of the same directory, say '/usr'. Instead they must be split into groups, say '/usr1' and '/usr2'; this is somewhat inconvenient, especially when space on one device runs out so that some users must be moved. The data movement can be done expeditiously, but the change in file names from '/usr1/...' to '/usr2/...' is annoying both to those people who must learn the new name and to programs which happen to have names built into them.

A final problem is not inherent to the basic design, but does aggravate the difficulty just mentioned. Disk block addresses for files are stored as 16-bit numbers, so a particular file system device can use no more than 65,526 blocks. This does not mean that the rest of a larger physical device is wasted, since there can be several logical devices per disk drive. But drives available today are an order of magnitude larger than 65K blocks, and chopping them up into eight or ten devices is unattractive.

I/O devices

Unix goes to some pains to efface differences between ordinary disk files and I/O devices such as typewriters, tapes, and line printers. An entry appears in the file system hierarchy for each supported device, so that the structure of device names is the same as that of file names. The same read and write system calls apply to devices as to disk files.

Moreover, the same protection mechanisms apply to devices as to files. Besides the traditionally available devices, device files exist for disk devices regarded as physical units outside the file system, and for absolutely addressed memory. The most important device in practice is the user's typewriter, since the standard input and output streams are attached to it. Because the typewriter channel is treated in the same way as any file (for example, the same I/O calls apply), it is easy to redirect the input and output of commands from the typewriter to another file, as explained in the next section.

Some differences are inevitable. For example, the system ordinarily treats typewriter input in units of lines, because character-erase and line-delete processing cannot be completed until a full line is typed. Thus if a program attempts to read some large number of bytes from a typewriter, it waits until a full line is typed, and then receives a notification that some smaller number of bytes has actually been read. All programs must be prepared for this eventuality in any case, because a read operation from any disk file will return fewer bytes than requested when the end of the file is encountered. Ordinarily, therefore, reads from the typewriter are fully compatible with reads from a disk file. A subtle problem can occur if a program reads several bytes, and on the basis of a line of text found therein calls another program to process the remainder of the input. Such a program works successfully when the input source is a typewriter, because the input is returned a line at a time, but when the source is an ordinary file the first program will have consumed input intended for the second. At the moment the simplest solution is for the first program to read one character at a time. A more general solution, not implemented, would allow a mode of reading wherein at most one line at a time was returned no matter what the input source.*

The user interface

The command interpreter, called the 'shell,' is the most important communication channel between the system and its users. The shell is not part of the operating system, and enjoys no special privileges. A part of the entry for each user in the password file read by the login procedure contains the name of the program to run initially, and for most users that program is the shell. This arrangement is by now commonplace in well-designed systems, but is by no means universal. Among its advantages are

*This suggestion may seem in conflict with our earlier disdain of 'records.' Not really, because it only affects the mode of reading information, not the way it is stored. The same bytes would be obtained in either read mode.

the ability to swap the shell even if the kernel is not swappable, so that the size of the shell is not of great concern. It is also easy to replace the shell with another program, either for testing purposes or to provide a non-standard interface.

The full language accepted by the shell is moderately complicated, because it performs a number of functions. Nevertheless the treatment of individual commands is quite simple and regular: a command is a sequence of words separated by white space (spaces and tabs). The first word is the name of the command, where a command is just an executable file. A full pathname, with '/' characters, may be used to specify the file unambiguously; otherwise an agreed-upon sequence of directories is searched. The only distinction possessed by a system-provided command is that it appears in a directory in the search path of most users. (A very few commands are build into the shell.) The other words making up a command line fall into three types:

- simple strings of characters
- A filename preceded by '<', '>', or '>>'
- A string containing a filename-expansion character.

The simple arguments are passed to the command as an array of strings, and thereafter are interpreted by that program. The fact that the arguments are parsed by the shell and passed as separate strings gives at least a start towards uniformity in the treatment of arguments; we have seen several systems in which arguments to various commands are separated sometimes by commas, sometimes by semicolons, and sometimes in parentheses; only a manual close at hand or a good memory tells which.

An argument beginning with '<' is taken to be a file which is to be opened by the shell and associated with the *standard input* of the command, which is the stream from which programs ordinarily read input; in the absence of '<' it is attached to the typewriter. Correspondingly, a file name beginning with '>' receives the standard output of commands; '>>' designates a variant in which the output is appended to the file instead of replacing it. For this mechanism to work it is necessary that I/O

to a typewriter be compatible with I/O to a file; the point here is that the redirection is specified in the shell language, in a convenient and natural notation, so that it is applicable uniformly and without exception to all commands. An argument specifying redirection is not passed to the command, which must go to some trouble even to discover if redirection has occurred. Other systems support I/O redirection (regrettably, too few), but we know of none with such a convenient notation.

An argument containing a file name expansion character is turned into a sequence of simple arguments which are the names of files. The character '*', for example, means 'any sequence of zero or more characters;' the argument '*.c' is expanded into a sequence of arguments which are the names of all files in the current directory whose names end with the characters '.c'. Other expansion characters specify an arbitrary single character in a file name or a range of characters (the digits, say).

Putting this expansion into the shell has several advantages: the code only appears once, so no space is wasted and commands in general need take no special action; the algorithm is certain to be applied uniformly. The only convention required of commands which treat files is to accept a sequence of file arguments even if the elementary action performed applies to only one at a time. For example the command which deletes a file could have been coded to accept only a single name, in which case argument expansion would be in vain; in fact it accepts a sequence of file names (however generated) and deletes all of them. Only occasionally is there any difficulty. For example, suppose the command 'save' transfers each of its argument files to off-line storage, so 'save *' would save everything; this works well. But the converse command 'restore' which might bring all the named arguments back on-line will not in general work as desired; 'restore *' would bring back only the files which currently exist (match the '*'), rather than all files saved.

One of the most important contributions of Unix to programming is the notion of *pipes* and especially the notation the shell provides for using them. A pipe is in effect

on the Honeywell had the same number of bytes (each of nine bits rather than eight) but there were 1280 bytes per block. With otherwise idle machines, the real times to accomplish the file copies were

<i>system</i>	<i>sec.</i>	<i>ms/hardware block</i>
Unix	25	26.0*
IAS	19	19.8
H6070	9	23.4

The machine on which the IAS test was run had newer disks, with essentially the same seek time, but with an average rotational latency of 8.3 ms. instead of 12.5 for ours. Also the time to transfer one block was .6 ms. instead of 1.9 for ours. If one corrects the Unix time for this hardware difference, the transfer rates become essentially the same, and the Honeywell rate is not far off when measured in blocks per second. No general statistical significance can be ascribed to this little experiment. Seek time, for example, must surely dominate the times measured, and there was in no case any attempt to optimize the placement of the input or output files. The results do seem to suggest, however, that the very flexible scheme for representing Unix files carries no great cost compared with at least two other systems.

The real time per block of I/O observed in the little test mentioned above was 26 milliseconds. Since the system overhead per block is 6 ms., most of which is overlapped, it would seem that the overall transfer rate of the copy might be nearly doubled if a block size of 1024 bytes were used instead of 512. There are some good arguments against this step. For example, space utilization on the disk would suffer noticeably; doubling the block size would increase the space occupied by files on our system by about 15 per cent, a number which increases in apparent importance when we consider that the free space is currently only 5 per cent. Increasing the

* For fair comparisons with the other systems the copies were timed on a particular file. The figure of 26 ms. per block is however quite close to the number obtained as an average for all I/O on the device over a period of hours or days.

block size would also force a decrease in the size of the system's buffer cache and lower its hit rate. However, this effect has not been reliably estimated.

Moreover, the copy program is an extreme case in that it is totally I/O bound, with no processing of the data. Most programs do at least look at the data as it goes by. Thus to sum the bytes in the file mentioned above requires 12 seconds of real time, five of which are user time spent looking at the bytes. To read the file and ignore it completely requires 11 seconds, with negligible user time. It may be concluded that the read-ahead strategy is almost perfectly effective, and that a program which spends as little as 50 microseconds per byte processing its data will not be significantly delayed waiting for I/O.

The basic system interface conceals physical aspects of file storage, such as blocks, tracks, and cylinders. Likewise the concept of a record is completely absent from the operating system proper and nearly so from the standard software. (By the term 'record' we mean an identifiable unit of information consisting either of a fixed number of bytes or of a count together with that number of bytes.) A text file, for example, is stored as a sequence of characters with new-line characters to delimit lines. This form of storage is not only efficient in space when compared with fixed-length records, or even records described by character counts, but is also the most convenient form of storage for the vast majority of text-processing programs, which almost invariably deal with character streams. Most important of all, however, is that there is only one representation of text files. One of the most valuable characteristics of Unix is the degree to which separate programs interact in useful ways; this interaction would be seriously impeded if there were a variety of representations of the same information.

We recall with a certain horrified fascination a system whose Fortran compiler demanded as input a file with variable-length records each of which was required to be 80 bytes long. (The prevalence of this sort of nonsense makes the following test of software flexibility (due to M. D. McIlroy) interesting to try when meeting new sys-

tems. It consists of writing a Fortran (or PL/I, or other language) program which copies itself to another file, running the program, and attempting to compile the resulting output. Most systems eventually pass, but usually only after an expert has been called in to mutter incantations which convert the data files generated by Fortran programs to the format expected by the Fortran compiler. In sum, we would consider it a grave imposition to require our users or ourselves, when mentioning a file, to specify the form in which it is stored.

For the reasons discussed above, Unix software does not use the traditional notion of 'record' in relation to files, particularly those containing textual information. But certainly there are applications in which the notion has use. A program, or self-contained set of programs, which generates intermediate files is entitled to use any form of data representation it considers useful. A program which maintains a large database in which it must frequently look up constant-sized entries given an index number may very well find it convenient to store the entries sequentially, in fixed-size units, sorted by index number; with some changes in the requirements, other file organizations become more appropriate. It is straightforward to implement any number of schemes within the Unix file system precisely because of the uniform, structureless nature of the underlying files; the standard software, however, does not include prewritten mechanisms to do it. As an example of what is possible, INGRES [Stonebraker et al.] is a relational database manager running under Unix which supports five different file organizations.

The structure of the filing system

On each file system device such as a disk the accessing information for files is arranged in an array starting at a known place. A file may thus be identified by its device and its index within the device. The internal name of a file is, however, never needed by users or their programs. There is a hierarchically arranged directory structure in which each directory contains a list of names (character strings) and the associated file index, referred implicitly to the same device as the directory. Since directories are


themselves files, the naming structure is potentially an arbitrary directed graph. Administrative rules restrict it to have the form of a tree, except that non-directory files may have several names (directory entries).


A file is named by a sequence of directories separated by '/' leading towards the leaves of the tree. The path specified by a name starting with '/' originates at the root; without an initial '/' the path starts at the current directory. Thus the simple name 'x' indicates the entry 'x' in the current directory; '/usr/dmr/x' searches the root for directory 'usr', searches it for directory 'dmr', and finally specifies 'x' in 'dmr'.

When the system is initialized, only one file system device is known (the 'root device'); its name is built into the system. More storage is attached by 'mounting' another device, which contains its own directory structure. When a device is mounted, the root of its directory tree is attached to a leaf of the currently accessible hierarchy. For example, suppose a device containing a subhierarchy is mounted on the file '/usr'. From then on, the original contents of /usr are hidden from view, and in names of the form '/usr/...' the '...' specifies a path from the root of the newly mounted device.

This file system design (described in detail in [Ritchie and Thompson]) is cheap to implement, is general enough to satisfy most demands, and has a number of virtues; for example, device self-consistency checks are straightforward. It does have a few peculiarities. For example, instantaneously-enforced space quotas, either for users or for directories, are relatively difficult to implement (it has been done, but is not in the standard system). Perhaps more serious, duplicate names for the same file ('links') while trivial to provide on a single device, do not work across devices; that is, a directory entry cannot point to a file on another device.

The implementation of mounted file systems is generally adequate, except for the problem, mentioned above, that links cannot be made across devices. Another limitation of the design is that an arbitrary subset of members of a given directory cannot be





**PWB/UNIX Papers
from the
Second International Conference on Software Engineering**

CONTENTS

An Introduction to the Programmer's Workbench	1
<i>T. A. Dolotta and J. R. Mashey</i>	
Using a Command Language as a High-Level Programming Language	6
<i>J. R. Mashey</i>	
Documentation Tools and Techniques	14
<i>J. R. Mashey and D. W. Smith</i>	
A User's Viewpoint on the Programmer's Workbench	19
<i>M. H. Bianchi and J. L. Wood</i>	



An Introduction to the Programmer's Workbench

T. A. Dolotta
J. R. Mashey

Bell Laboratories
Piscataway, New Jersey 08854

Keywords: Program development facility, UNIX.

Abstract: The Programmer's Workbench (PWB) is a specialized computing facility dedicated to satisfying the needs of developers of computer programs. The PWB might well be called a "human-end" computer; like "front-end" and "back-end" computers, it improves productivity by efficient specialization. It provides a convenient working environment and a uniform set of programming tools to a diverse group of programming projects. These projects produce software for various "target" computers, including IBM System/370 and UNIVAC 1100 systems of much greater size than the PWB machines. The projects range in size from several people up to several hundred. The first PWB machine was installed in October, 1973; usage, acceptance, and interest have grown rapidly since that time. The PWB currently supports about 110 time-sharing terminals, utilizing a network of four DEC PDP-11 computers, all running the UNIX Time-Sharing System. The PWB adds tools to UNIX to support large projects. This paper gives an overview of the PWB and its development; further details appear in the five following companion papers [IBIA76A, IXL76B, KNU76A, MAS76A, MAS76B].

1. INTRODUCTION

The Programmer's Workbench (PWB) is a specialized computing facility dedicated to supporting large software development projects. Although it performs many of the functions that have been suggested for projects such as the Program Support Library [LUP74A], the Development Support Library [BAK75A], the Automated Project Management Information System [BRA75B], or the Software Factory [BRA75A], some of its goals are quite different, and its implementation is even more so. Each of these systems may be categorized as a program development facility (PDF), i.e., a system or package of programs meant specifically to support software development.

The PWB is a production system that has been used for several years in the Business Information Systems area of Bell Laboratories. It supports a user community of about 700 people, and can handle 110 simultaneous on-line users. Although the PWB is still growing and evolving at a rapid rate, accumulated experience provides strong support for the viability, economy, and effectiveness of its approach.

This paper explains the rationale for the existence of the PWB and describes its history and current status. It also serves to supply the background for five companion papers, which, in turn, supply examples and evidence that support the assertions and opinions presented here.

2. THE PWB CONCEPT

The PWB concept embodies the following ideas:

- Program development and execution of the resulting programs are two radically different functions. Much can be gained by assigning each function to a computer best suited for it. Thus, as much of the development as possible should be done on a computer dedicated to that task, while execution should occur on another computer, called a "host" or "target" system.
- Although there may be several target systems, possibly supplied by different vendors, the PDF should present a single,

uniform interface to its users. Existing PWB targets include IBM System/370 and UNIVAC 1100 computers; in some sense, the PWB is a target also, in that it is built and maintained with its own tools.

- A PDF should be implemented on several computers of moderate size, even when the target machines consist of very large systems. The rationale for this idea is given in sections 4.1 and 4.6. The PWB currently is implemented as a network of four large DEC PDP-11's (three 11/45's and one 11/70), all operating under the UNIX Time-Sharing System [RIT74A].

Although the PWB is a special-purpose system in the same sense as a "front-end" or "back-end" [CAN74A] computer, its goal is to be a "human-end" computer. As shown in Figure 1, it provides the primary interface between program developers and their target computer(s). Unlike a typical "front-end," the PWB supplies a separate, uniform environment in which people perform their work.




Figure 1. PWB Interface for Users

3. DISADVANTAGES—REAL AND IMAGINARY

The PWB approach of using small computers to build a PDF for use with much larger targets has both good and bad points. Fortunately, many of the potential disadvantages that caused concern at the outset of the PWB project have not materialized, and the actual disadvantages are far outweighed by the benefits [IV175A]. The following points apply mainly to our implementation of the PWB; some may have little relevance for other PDFs.

3.1 Additional Hardware Costs

The first disadvantage is that of requiring more computer hardware and communications gear. This is a real disadvantage only if one has target systems with excess capacity. In our case, the reverse was true: overall hardware costs were probably decreased by using several minicomputers instead of obtaining additional large systems or upgrading those already installed. An associated (and real) problem is that of the overhead of separate contracts, maintenance, training, etc.

3.2 Increased Complexity of Total System

The second disadvantage is the increased number of nodes and links in the resulting network. There is a potential for disaster if failure of a few components can make the rest unusable. This can (and has sometimes been) a real problem, but is lessened, if not completely eliminated, by providing appropriate redundancy, and by configuring the overall system so that individual components can perform useful stand-alone work. As an example, when the PWB is up, users can edit programs and queue jobs for later transmittal to a target, even if that target or the communications links to it are down.

3.3 Data and Function Splitting

The third problem is that of splitting data and capabilities between machines. Communication links must be fast enough to avoid unacceptable delays. Fortunately, it turns out that many types of data need exist only in one place. For example, the PWB concentrates on the handling of text files containing source programs, control statements, and documentation. A target only keeps such things temporarily, but it does keep its own object programs, which are seldom, if ever, seen by the PWB.

3.4 Incompatibilities

Incompatible character sets and the resulting data conversion costs are the fourth problem. In practice, this has been at most a minor irritation. It may not be a problem at all if target and PDF computers are chosen to use the same character sets.

3.5 Load Balancing

The fifth disadvantage is the possible loss of flexibility in load balancing, because some resources are dedicated to specific tasks and cannot be used for others.

3.6 Limits on Size and Speed

Large systems are capable of many actions far outside the range of small ones. For example, as currently implemented, the PWB does not offer interactive debugging on the target systems.

3.7 Duplication of Software

The final problem is that of the duplication of software, not only between the target and PWB, but between separate PWB processors. For example, each PWB CPU has an independent copy of the system software. However, this software is relatively small, and consumes much less space than user files and programs.

Another aspect of the problem is the need to distribute and coordinate the installation of software on separate systems. One compensation is the ability to test new or modified programs (including the operating system) on one system before installing it on all the others.

4. ADVANTAGES

The PWB approach is no panacea for many of the problems facing the software industry today. However, it appears to be very effective in the presence of certain conditions, each making the PWB approach advantageous. It is hardly accidental that the PWB originated at an installation where *all* these conditions exist.

4.1 Gain by Effective Specialization

The computer requirements of software *developers* often diverge quite sharply from those of the *users* of that software. This observation seems relevant to many kinds of software projects, but is especially true in the production of large, interactive, data-base-oriented systems. In our opinion, the following are some of the primary needs of software developers:

- Convenient, inexpensive, and continually available interactive computing services.
- A file structure oriented to interactive use, as opposed to one grafted as an additional layer onto a batch-oriented structure. The overall structure and individual files must be quickly changeable, should avoid unnecessary involvement with hardware details, and should *never* require explicit allocation of storage. This observation derives from watching program developers waste much effort on managing disk space, rather than doing their work.
- Powerful, human-engineered, and surprise-free tools for scanning and editing text files. Much programming activity consists of manipulation of text.
- A flexible and easy-to-use command language. It should be a programming language in its own right, but still simple enough to allow quick learning.
- Extensive document preparation facilities.

- Facilities to help solve *small* data base management problems quickly and cheaply.
- Adaptability to rapid organizational and personnel changes.
- Guaranteed service during normal working hours. PWB users can survive more easily the temporary loss of a target machine or communication links than the loss of the PWB itself.

On the other hand, users of the end product may have any or all of the following needs:

- Hardware large and fast enough to run the end product, possibly under stringent real-time and deadline constraints.
- File structure and access methods that can be optimized to handle large amounts of data.
- Transaction-oriented teleprocessing facilities.
- The use of a specific type of computer and operating system, in order to meet any one of a number of possible, externally-imposed requirements.

Although, in principle, it is possible to conceive of a system that satisfies both of the above sets of requirements in a cost-effective way, we are not aware of the existence of such a system. If it is necessary to choose one set over the other, the developers often lose the battle because the end users are usually paying for the system. Given the fact that software costs already exceed hardware costs, and are expected to do so to an even greater extent in the future [BOE73A, DOL76C], this may be a very unwise approach in the long run.

Given the prevailing costs and natures of large and small computers, it would appear that small ones are much more suited to the usual needs of program developers. Furthermore, assuming that a PDF should have duplicated systems in order to guarantee service, it seems logical to duplicate small, cheap computers instead of large, expensive ones.

4.2 Multi-Vendor Installations

It is desirable to have a uniform set of tools in order to ease training and to permit the transferral of personnel between various projects. The creation of such a set of tools is made difficult by the differences in file structures, command languages, and communications protocols of the various targets. These differences are very troublesome; they are qualitatively different from problems encountered in transferring, say, COBOL programs between various target systems. Nevertheless, they must be dealt with. As a result, a few PWB tools are target-specific, but, wherever possible, they possess target-independent user interfaces.

4.3 Changing Environments

Changes to hardware and software occur and cause problems even in single-vendor installations. Such changes may be disastrous if they affect both development and production environments at the same time. This problem is at least partially solved by using a separate program development system. The approach also helps programmers "get going" on a project before the production hardware becomes available.

As an example, over the last year, there have been major reconfigurations in both the hardware and software of all of our target systems, and in the geographic work locations of all of our users. Because of the "insulation" that PWB provides, most of our users were relatively unaffected by these changes.

4.4 Effective Testing of Terminal-Oriented Systems

It is difficult enough to test small batch programs; effective testing of large, interactive, data base management systems is far more difficult. It is especially difficult to perform load testing when the same computer is both generating the load and running the application program being tested. It is much simpler and more realistic to perform testing with the aid of a separate computer.

4.5 Availability of Better Software

Many time-sharing systems for large computers often retain significant vestiges of batch processing. Of necessity, some of them have evolved over a long period of time, and may contain design elements that would not exist if they were redone in the light of current knowledge. Separation of support functions onto an appropriate minicomputer may offer an easy way to gain access to more up-to-date software. For example, much of the stimulus for the PWB arose from the desire to make use of the simple, elegant, and powerful UNIX Time-Sharing System.

4.6 Human Orientation and Sizing

In many cases, operating systems, especially interactive ones, have become incredibly complex and incomprehensible. Small systems can be made more comprehensible, friendly, and adaptive, and an individual user can have more impact on their development. Many people welcome *any* movement in the direction of simplicity: "Small is Beautiful" applies to this area, as well as many others [SCH73B].

5. CURRENT STATUS

This section represents a snapshot of the status of the PWB as of June, 1976. Due to the continuing, rapid growth in its usage, this snapshot will soon to be obsolete.

5.1 Hardware Configuration

Figure 2 displays the current structure of the PWB and of its target computers.

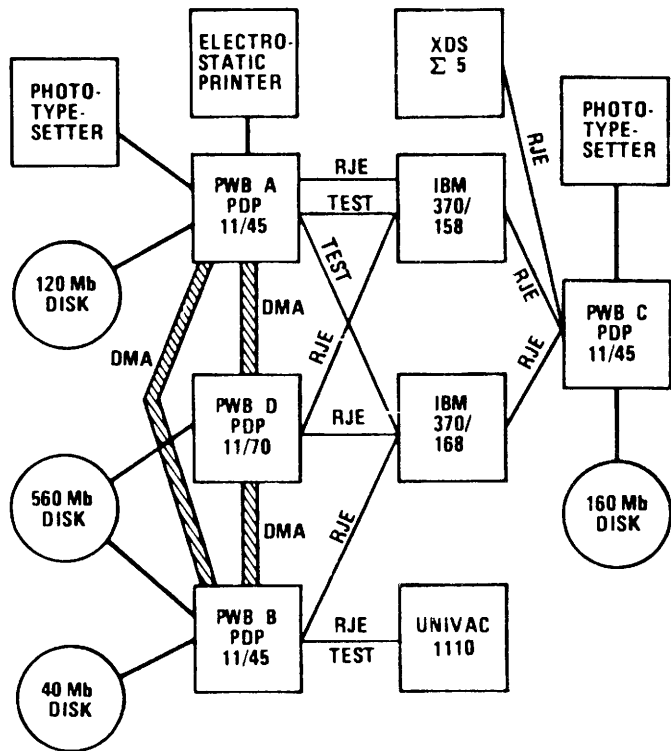


Figure 2. PWB Hardware Configuration

Several notes are in order regarding Figure 2:

- Systems A, B, and D are all interconnected via direct memory-to-memory (DMA) links.
- The communications links between PWB and targets consist of a number of 48kb, 9.6kb, and 48kb channels; these links are actually implemented as sub-channels of several multiplexed 230kb lines.
- Each PDP-11 has 124K words of memory (248K bytes). Systems A, B, and D are located in one building (Piscataway, N.J.), the IBM and UNIVAC systems are in another

building a few miles away; System C is located about 20 miles away (at Murray Hill, N. J.), as is the XDS SIGMA 5.

- System A is used primarily by the people who develop and support the PWB, by users of some testing facilities (DOL76B), and by users of the phototypesetter.
- System B is primarily used by the developers of a very large, UNIVAC-based system, and also supports additional documentation activities.
- System C supports a large IBM-based project, as well as some members of a project using XDS and other computers.
- System D (a PDP 11/70) is used by developers of other IBM-based systems. Some additional text processing (typing pool) is also done here.
- Systems B and D share disk drives and each can easily be used as backup for the other.
- Temporarily, System C has no convenient backup. This situation will be rectified in the near future. In the next upgrade, Systems A and B will become PDP 11/70's, and a PDP 11/45 (System E) will be moved to join System C.
- Equivalent monthly lease & maintenance costs total approximately \$26,000 for these four systems, exclusive of terminals and communications lines.

5.2 Sample Usage Data

Figure 3 displays data obtained at the end of June, 1976.

System	Login Names	Files	Directories	Dial Ports	Prime Connect Hours/Day
A	134	16266	1334	16	71
B	160	27924	1828	26	124
C	156	10430	872	20	121
D	224	32254	2016	48	312
Totals	674	86874	6050	110	628

Figure 3. PWB Usage

In many cases, the same login name is used by many people who work together. In a few cases, a single person possesses several such names. Some login names exist for training new users and for experimental use, and are used only occasionally.

The connect times are for 8-hour prime shifts only and represent averages computed from one month's totals. Of the 880 possible prime connect hours, 71% were actually consumed. It is difficult to raise this percentage significantly; during some parts of the day, *all* ports are in use. The prime shift accounts for approximately 75% of all connect hours. Usage during other shifts includes stand-alone and other testing, maintenance procedures, and execution of some extremely time-consuming user programs that need no human interaction.

6. CURRENT FACILITIES

6.1 UNIX

Much of the impetus for the development of the PWB arose from the desire to utilize the UNIX environment. Although the basic flavor of UNIX has been maintained as much as possible within the PWB, extensions have been necessary to make it more suitable as the operating system for the PWB. The PWB is currently the largest known UNIX installation, in terms of number of systems, disk space, and number of users: some changes were needed to support such an atypically large user community. In particular, much effort has gone into improving the robustness and efficiency of the operating system itself. It is a tribute to the clarity and adaptability of UNIX that it can be easily adjusted to a wide range of configurations.

Various UNIX commands were changed in significant ways. The editor was modified to better support users dealing with fixed-format source programs and data. Many additions and

modifications were made to increase the performance of the command interpreter, improve it for use by multi-person teams, and upgrade its capabilities as a programming language [MAS76A].

6.2 Remote Job Entry (RJE) Facility

The Remote Job Entry (RJE) facility includes the following:

- A command used to generate job streams for target systems: it performs nested file inclusions, keyword substitutions, prompting, and character translation (e.g., ASCII to EBCDIC). It also includes a generalized interface to other PWB commands, so that all or parts of job streams can be generated dynamically from the output of other UNIX commands.
- Transmission subsystems used to handle all communications activity: these are target-specific, but are not visible to the users.
- Status reporting: users may inquire about the status of jobs on the target systems, and can elect to be notified in various ways (e.g., on-line or in absentia) of the occurrence of major events in the processing of these jobs.
- Output retrieval: users may route the target's output to remote printers, or may elect to have all or part of it returned into PWB files. These files can be examined by various scanners, or manipulated using many different commands.
- Supporting tools: some users need to create their own job streams. Others want only to transfer files among systems or print PWB files on target system printers. Various utility programs exist to simplify these functions. Other utilities are used to aid in converting existing target source programs to formats more appropriate for the PWB. In particular, sequence numbers and trailing blanks are eliminated, and other strings of blanks are converted to tabs if possible.

[BIA76A] includes a brief description and examples of RJE usage as seen by the users.

6.3 Source Code Control System (SCCS)

The Source Code Control System permits unusually powerful control over changes to *modules* (i.e., files of text—source code, documentation, data, etc.). It records every change made to a module, can recreate a module as it existed at any point in time, controls and manages any number of concurrently existing versions of a module, and offers many useful audit and administrative features. SCCS is described in [ROC75A].

6.4 Modification Request Control System (MRCS)

Any extensive software project usually evolves mechanisms for requesting changes, reporting errors, and tracking the progress of modifications. Many projects have created systems for handling these items, but the systems are often specific to one project. MRCS is a generalized system intended for use by many projects, and it is described in [KNU76A].

6.5 Document Preparation

UNIX has traditionally provided good documentation tools. Additional commands have been added, and the existing tools put to use in various novel ways. The PWB currently supports a great deal of documentation work, for programmer and non-programmer alike. Many of the documentation efforts are difficult to distinguish from data base applications. Descriptions for the tools and techniques of document preparation are given in [MAS76B].

6.6 Test Drivers

The PWB is often used to run various kinds of tests of IBM and UNIVAC data base management systems and of data base applications implemented on these systems. The PWB contains two test drivers that can generate repeatable tests for very complex interactive systems; these drivers are used both to measure performance under well-controlled load and to help verify the initial and continuing correct operation of these systems and applications while they are being built and modified. One driver

simulates a Teletype® cluster controller of up to four terminals, and is used to test programs running on UNIVAC 1100 series machines. The other driver (LEAP-[DOL76B]) simulates an IBM 3270 cluster controller managing up to eight terminals.

7. HISTORY AND DESIGN PHILOSOPHY

7.1 Brief History

The PWB concept was suggested in April, 1973, and the first PDP 11/45 was installed in October, 1973. This machine was used at first by the development department for its own education and experimentation, while the original RJE, SCCS, and LEAP components were constructed. Additional systems were installed in October, 1974 and in May, 1975. The PDP-11/70 arrived in October, 1975; further upgrades are planned for 1976.

At first, the PWB was an experimental project that faced considerable indifference from a user community heavily oriented to large computer systems, working under difficult schedules, and a bit wary of what then seemed like a radical idea. However, as word spread about the system, demand for service began to outrun our ability to supply it. When forced to make six-month usage forecasts, users consistently underestimated the extent of their usage, because they kept discovering unexpected applications for PWB facilities. New hardware was (and is) continually being acquired to meet the demand.

Several large projects are awaiting the arrival of additional hardware before they can completely convert to the PWB. We expect the growth of PWB facilities at Bell Laboratories to continue in the foreseeable future.

7.2 Design Philosophy

Early in the PWB development cycle, many spirited discussions were held regarding the nature of the overall design approach to be taken. One proposed approach was that of first designing the PWB as a completely integrated facility, then implementing it, and then obtaining users for it. A much different approach has actually been followed. Its elements are:

- Get users on the system quickly, and let their needs and problems drive the design.
- Build software quickly, and expect to throw much of it away.
- Build many small, independent programs, rather than large, interrelated ones. This fits well with the nature of UNIX.
- Use as few different file formats as possible, and keep them simple. This avoids the need for a plethora of "utilities" needed to deal with a wide variety of formats, and maximizes the use of existing commands.
- Monitor user problems and rearrange functions or add new ones as needed.
- Design each new feature to be as consistent with existing ones as possible, in order to maintain an environment that is simple, coherent, and conducive to productive use.
- When incompatible changes are necessary, let them derive from user demands, or make sure they offer so much benefit that users will be glad to have them.

This approach may appear chaotic, but evidence to support its desirability can be found in [BRO75A, p. 116] and [NAU69A, pp. 19, 22, 32, 40, 41, 47, 73, 95], for example. In practice, it seems to work better than designing "perfect" systems that turn out to be obsolete or unusable by the time they are implemented. Of course, we are lucky in being able to utilize an operating system that both permits and encourages this kind of approach.

8. CONCLUSIONS

We have given a brief summary of the rationale for the PWB, its advantages and disadvantages, its current status, and the philosophy behind the approach taken. No claim is made that the PWB can solve everyone's programming problems. We do claim that the PWB approach is quite appropriate for many problems.

existing today, that the directions of change in the computer industry (DOL76C) may make it even more applicable in the future, and that it does handle many real problems, as evidenced by considerable live experience.

ACKNOWLEDGEMENTS

The PWB concept was first suggested by E. L. Ivie (IV175A). The authors of the five companion papers, as well as many of our other colleagues, have contributed to the design, implementation, and continuing improvement of the PWB. UNIX itself, without which PWB could not have been built, was developed by members of the Bell Laboratories' Computing Science Research Center. Finally, much of the success of the PWB can be attributed to a user population willing to try new things and to give us the feedback necessary to make these things useful.

REFERENCES

All references for this and the five companion papers [BIA76A, DOL76B, KNU76A, MAS76A, MAS76B] are given below.

- ADR73A Applied Data Research. *The LIBRARIAN: User Reference Manual*. Report P112L, Applied Data Research, Inc., Princeton, N. J., 1973.
- BAK75A Baker, F. T. Structured Programming in a Production Programming Environment. *Proc. Int. Conf. on Reliable Software*, April 21-23, 1975, 172-85.
- BIA76A Bianchi, M. H., and Wood, J. L. A User's Viewpoint on the Programmer's Workbench. *Proc. Second Int. Conf. on Software Engineering*, Oct. 13-15, 1976.
- BOE73A Boehm, B. W. Software and its Impact: A Quantitative Assessment. *Datamation* 19, 5 (May 1973), 48-59.
- BRA75A Bratman, H., and Court, T. The Software Factory. *Computer* 8, 5 (May 1975), 28-37.
- BRA75B Bratman, H. Automated Techniques for Project Management and Control. In *Practical Strategies for Developing Large Software Systems*, ed. E. Horowitz, Addison-Wesley, Reading, Mass., 1975, 193-211.
- BR075A Brooks, F. P., Jr. *The Mythical Man-Month*. Addison-Wesley, Reading, Mass., 1975.
- BRU76A Brunt, R. F., and Tuffs, D. E. A User-Oriented Approach to Control Languages. *Software—Practice and Experience* 6, 1 (Jan. 1976), 93-108.
- CAN74A Canada, R. H., Harrison, R. D., Ivie, E. L., Ryder, J. L., and Wehr, L. A. A Back-End Computer for Data Base Management. *Comm. ACM* 17, 10 (Oct. 1974), 575-82.
- COL76A Colijn, A. W. Experiments with the KRONOS Control Language. *Software—Practice and Experience* 6, 1 (Jan. 1976), 133-36.
- COM76A COMTEN HyperFASTER Concepts and Facilities Manual. COMTEN, Inc., 2 Research Court, Rockville, Md.
- COW75A Cowan, R. M. Burroughs B6700/B7700 Work Flow Language. In *Command Languages—Proc. IFIP Working Conference on Command Languages*, ed. C. Unger, North Holland, Amsterdam, 1975, 153-66.
- DEC74A Digital Equipment Corp. Option Description, DQS11-A/B Communications Controller. Form CSS-MU-F-32-3A, Digital Equipment Corp., Maynard, Mass., 1974.
- DOL69A Dolotta, T. A., and Irvine, C. A. Proposal for a Time Sharing Command Structure. In *Information Processing 68—Proc. IFIP Congress 1968*, vol. 1, North Holland, Amsterdam, 1969, 493-98.
- DOL76A Dolotta, T. A., and Mashey, J. R. An Introduction to the Programmer's Workbench. *Proc. Second Int. Conf. on Software Engineering*, Oct. 13-15, 1976.
- DOL76B Dolotta, T. A., Licwinko, J. S., Menninger, R. E., and Roome, W. D. The LEAP Load and Test Driver. *Proc. Second Int. Conf. on Software Engineering*, Oct. 13-15, 1976.
- DOL76C Dolotta, T. A., et al. *Data Processing in 1980-1985*. Wiley-Interscience, New York, 1976.
- GRI69A Greenbaum, H. J. *A Simulator of Multiple Interactive Users to Drive a Time-Shared Computer System*. MIT, Cambridge, Mass., 1969 (avail. from NTIS, AD 686 988).
- IBM72A IBM OS/VS Utilities. Form GC35-0005, IBM Corp., White Plains, N. Y., 1972.
- IBM73A IBM Introduction to Programming the IBM 3270. Form GC 27 6999, IBM Corp., White Plains, N. Y., 1973.
- IBM73B IBM IBM 3270 Information Display System Component Description. Form GA21 2349, IBM Corp., White Plains, N. Y., 1973.
- IBM74A IBM IMS/VS System Operators Reference Manual. Form SH20-9028, IBM Corp., White Plains, N. Y., 1974.
- IBM75A IBM IMS/VS System Programming Reference Manual. Form SH20-9027, IBM Corp., White Plains, N. Y., 1975.
- IBM75B IBM IMS/VS General Information Manual. Form GH20-1260, IBM Corp., White Plains, N. Y., 1975.
- IBM75C IBM DB/DC Driver System General Information Manual. Form GH20-1639, IBM Corp., White Plains, N. Y., 1975.
- IBM76A IBM Teleprocessing Network Simulator (TPNS) Program Reference Manual. Form SH20-1823, IBM Corp., White Plains, N. Y., 1976.
- IV175A Ivie, E. L. *The Programmer's Workbench—A Machine for Software Development*. Unpublished Report, Bell Laboratories, May 19, 1975.
- JAM75A James, D. L., and Lambert, D. W. *Remote-Terminal Emulator (Design Verification Model)—Introduction and Summary*. Mitre Corp., Bedford, Mass., 1975 (avail. from NTIS: AD A007 827).
- KER75A Kernighan, B. W., and Plauger, P. J. Software Tools. *Proc. First National Conference on Software Engineering*, Sept. 11-12, 1975, 8-13.
- KER75B Kernighan, B. W., and Cherry, L. L. A System for Typesetting Mathematics. *Comm. ACM* 18, 3 (Mar. 1975), 151-56.
- KER76A Kernighan, B. W., and Plauger, P. J. *Software Tools*. Addison-Wesley, Reading, Mass., 1976.
- KNU76A Knudsen, D. B., Barofsky, A., and Satz, L. R. A Modification Request Control System. *Proc. Second Int. Conf. on Software Engineering*, Oct. 13-15, 1976.
- LUP74A Luppino, F. M., and Smith, R. L. *Programming Support Library (PSL) Functional Requirements*. Structured Programming Series, vol. 5, IBM Federal Systems Div., 1974 (avail. from NTIS, AD A003 339).
- MAS76A Mashey, J. R. Using a Command Language as a High-Level Programming Language. *Proc. Second Int. Conf. on Software Engineering*, Oct. 13-15, 1976.
- MAS76B Mashey, J. R., and Smith, D. W. Documentation Tools and Techniques. *Proc. Second Int. Conf. on Software Engineering*, Oct. 13-15, 1976.
- NAU69A Naur, P., and Randell, B., eds. *Software Engineering*. Scientific Affairs Division, NATO, Brussels 39, Belgium, 1969.
- OSS74B Ossanna, J. F. NROFF User's Manual, 2nd ed. Unpublished Report, Bell Laboratories, September 11, 1974.
- PUL68A Pullen, E. W., and Shuttee, D. F. MUSE: A Tool for Testing and Debugging a Multi-Terminal Programming System. *Proc. 4/FIPS Spring Joint Computer Conference*, vol. 32, 1968, 491-502.
- RIT74A Ritchie, D. M., and Thompson, K. The UNIX Time-Sharing System. *Comm. ACM* 17, 7 (July 1974), 365-75.
- RIT75A Ritchie, D. M. C Reference Manual. Unpublished Report, Bell Laboratories, February 1, 1975.
- ROC75A Rochkind, M. J. The Source Code Control System. *IEEE Trans. on Software Engineering*, SE-1, 4 (Dec. 1975), 364-70.
- SCH73B Schumacher, E. F. *Small is Beautiful*. Harper & Row, New York, 1973.
- SIM74A Simpson, D., ed. *Job Control Languages—Past, Present, and Future*. National Computing Centre Ltd., Quay House, Quay St., Manchester, England, 1974.
- THO75A Thompson, K. The UNIX Command Language. In *Structured Programming—International Computer State of the Art Report*, Infotech Information Ltd., Maidenhead, Berkshire, England, 1975, 375-84.
- THO75B Thompson, K., and Ritchie, D. M. UNIX Programmer's Manual 6th ed. Unpublished Report, Bell Laboratories, May, 1975.
- UNG75A Unger, C., ed. *Command Languages—Proc. IFIP Working Conference on Command Languages*. North-Holland, Amsterdam, 1975.

Using a Command Language as a High-Level Programming Language

J. R. Mashey

Bell Laboratories
Piscataway, New Jersey 08854

Keywords: Command languages, command interpreters, UNIX.

Abstract: The command language for the Programmer's Workbench (PWB) utilizes an extended version of the standard UNIX shell program, plus commands designed mainly for use within *shell procedures* (command files). Modifications have been aimed at improving the use of the shell by large programming groups, and making it even more convenient to use as a high-level programming language. In line with the philosophy of much existing UNIX software, an attempt has been made to add new features only when they are shown necessary by actual user experience in order to avoid contaminating a compact, elegant system through "creeping featurism." By utilizing the shell as a programming language, PWB users have been able to eliminate a great deal of the programming drudgery that often accompanies a large project. Many manual procedures have been quickly, cheaply, and conveniently automated. Because it is so easy to create and use shell procedures, each separate project has tended to customize the general PWB environment into one tailored to its own requirements, organizational structure, and terminology. A summary is given of the usage patterns revealed by a survey of 1,725 existing shell procedures.

1. INTRODUCTION

A good operating system command language (CL) is an invaluable tool for large programming projects. A common occurrence in such projects is the diversion of significant resources from building the end product to creating internal support programs and procedures. If a CL is a flexible programming language, it can be used to solve many internal support problems, without requiring compilable programs to be written, debugged, and maintained. Its most important advantage is the ability to do the job *now*.

The Programmer's Workbench (PWB) environment [IVI75A, IXI76A] supports projects ranging in size from several people to those involving hundreds. The PWB CL is a modified version of the UNIX *shell* language [RIT74A, THO75A]. Experience with it has shown it to be an effective tool for automating procedures and eliminating programming drudgery. Nearly two thousand CL procedures have been inspected to observe real-life usage. Many of the assertions made in this paper are based on the results of this survey.

Existing CLs include a wide range of structures and abilities, ranging from those quite similar to assembly language to a few having the appearance and abilities of high-level programming languages [BRU76A, COW75A]. A diversity of viewpoints exists regarding the relative importance of various CL characteristics [DOU69A, SIM74A, UNG75A]. Several years of PWB experience suggests that the following qualities are valuable in a CL:

- *Ease of on-line use*—Many advantages can be gained by replacing keypunches with terminals, even when jobs are prepared for batch execution. To support this approach, CL syntax should emphasize simplicity and avoid the need for redundant typing.
- *Convenient use of the same language as an on-line CL and as a programming language*—It must be possible to build and store sequences of commands (*CL procedures*) that can be invoked at a later time.

- *Interchangeability of programs and CL procedures*—No user should need to know whether a given command is implemented as a compiled, executable program or as a CL procedure. Syntactic differences should be avoided.
- *Ease of creation and maintenance of CL procedures*—It should require very little effort to write a CL procedure, save it for later use, and maintain it. The overhead of writing a small procedure should be especially small; as will be seen later, many procedures are only a few lines long.
- *Organization of CL procedures*—It must be possible to share easily the use of procedures among people in a way consistent with organizational structure.
- *Programming language features*—It must be possible to write procedures with convenient conditional branching, looping, argument handling, variables, string manipulation, and occasional arithmetic. Thus, many of the capabilities of a typical procedural language are *necessary* for a good CL. However they may not be *sufficient*: a CL may need additional capabilities, a somewhat different syntax, and a radically different set of priorities concerning the importance of various constructs. For example, a crucial CL capability is that of connecting executing programs in a variety of ways. Pattern-matching and other string manipulation operations are quite useful. Arithmetic operations are also helpful, but seem to be much less important than the others.
- *Separation from operating system*—A CL interpreter should normally be an ordinary command, treated like other commands and definitely *not* embedded in the heart of the operating system. The operating system should *not* be viewed as an implementation of the CL, but as an environment that can support a variety of good CLs. Different users should be able to have different CLs if they feel like it. This approach permits experimentation and evolution without bothering other users. It often leads to a "survival-of-the-fittest" behavior. Mutation occurs, lives, and dies on their merits, and eventually breeds hybrids containing features found useful from actual practice.

2. OVERVIEW OF THE UNIX ENVIRONMENT

Full understanding of some later discussions depends on familiarity with UNIX. [RIT74A] is a definite prerequisite, and it would be helpful to read at least one of [KER75A, KER76A, THO75A]. For completeness, a short overview of the most relevant concepts is given below.

2.1 File System

The UNIX file system's overall structure is that of a rooted tree composed of *directories* and other files. A *file name* is a sequence of characters. A *path name* is a sequence of directory names followed by a file name, each separated from the previous one by slash ("/"). If a path name begins with a "/", the search for the file begins at the root of the entire tree; otherwise, it begins at the user's *current directory*. (The first type of name is often called an *absolute path name* because it is invariant with regard to the user's current directory.) The user may change current directories at any time by using the *chdir* command. In general, file names and path names can be used in the same ways. Some sample names are:

/	root of the entire file structure.
/bin	directory of commonly used public commands.
/u0/tnds/tf/jtb/bin	a path name typical of multi-person programming projects. This one is a private directory of commands belonging to person "jtb" of group "tf" in project "tnds".
bin/umail	a name depending on the current directory: it refers to file "umail" found in subdirectory "bin" of the current directory. If the current directory is "/", it names "/bin/umail". If the current directory is "/u0/tnds/tf/jtb", it names "/u0/tnds/tf/jtb/bin/umail".

Large projects require the ability to quickly and easily modify directory structures to fit changing needs. In particular, the "current directory" feature makes it possible for each person to move around in the file system and work where most convenient. This allows simple names to be used, even when the current directory is many levels deep in the structure. It also permits individual directories to remain fairly small, lessening the load on both human and computer; i.e., "locality of reference" is good for the performance of *both*.

2.2 Processes and Their Interactions

An *image* is a computer execution environment, including core image, register values, current directory, status of open files, information recorded at login time, and various other items. A *process* is the execution of an image; most UNIX commands execute as separate processes. One process may spawn another using the *fork* system call, which duplicates the image of the original (*parent*) process. The new (*child*) process may continue execution of the image, or may abandon it by issuing an *exec* system call, which initiates execution of another program.

Processes use independent address spaces and data segments,¹ and communicate in a limited number of ways:

- *Open files*—a child inherits the parent's open files, and can manipulate the associated read/write pointers thus shared with the parent. This ability permits processes to share the use of a common input stream in various useful ways. In particular, an open file possesses a *pointer* that indicates a location in the file, and is modified by various operations. *Read* and *write* copy a requested number of bytes from (to) a file, beginning at the location given by the current value of the pointer. As a side effect, the pointer is incremented by the number of bytes transferred, yielding the effect of sequential I/O. *Seek* can be used to obtain random-access I/O; it sets the pointer to an absolute location within the file, or to a location offset from the end of the file or the current pointer location.
- *Arguments*—a sequence of arguments (character strings) can be passed from one program to another via *exec*.
- *Return code*—when a process terminates, it can set a numeric return code that is available to the process's parent.
- *Files*—some programs arrange conventions to share files in various ways, or to use files of specified names.
- *Pipes*—pipes are interprocess channels that are similar to files in ways of access, but allow very convenient handling of the "producer-consumer" relationship between programs executing in parallel. The "producer" writes into one end of a pipe, while the "consumer" empties it by reading from the other end. Because UNIX handles details of buffering and synchronization, neither program needs explicit information about the other's activities.

¹ The text segment of a reentrant program is shared by all processes executing that program. Almost all programs are reentrant.

Limiting interprocess communication to a small number of well-defined methods is a great aid to uniformity, understandability, and reliability of programs. It encourages the packaging of functions into small programs that are easily connected. The pipe mechanism is especially desirable, both for human comprehension and for computer performance [THO75A, KER75A, KER76A].

3. SHELL BASICS

Most UNIX users utilize the CL provided by a program called the *shell*. It reads input from a terminal or file and arranges for the execution of the requested commands. The shell is a small program (about 20 pages of C code); many CL functions are actually supported by independent programs that work with the shell, but are not built into it. The discussion is adapted from [THO75A, THO75B].

3.1 Commands

A *command* is a sequence of non-blank arguments separated by blanks. The first argument specifies the name of the command to be executed; the remaining items are passed as arguments to the command executed. The following line requests the *pr* command to print files a, b, and c:

```
pr a b c
```

If the first argument names a file that is marked as *executable*² and is actually a load module, the shell (as parent) spawns a new (child) process that immediately executes that program. If the file is marked executable, but is neither a load module nor a directory, it is assumed to be a *command file* (*shell procedure*). A command file is a file of ordinary text—shell command lines and possibly lines to be read by other programs. In this case, the shell spawns a new instance of itself to read the file and execute the commands included in it.

From the user's viewpoint, executable programs and shell procedures are invoked in exactly the same way. The shell determines which implementation has been used, rather than requiring the user to do so. Most operating systems can execute existing load modules without requiring the user to state the source language used to produce the load module. CL procedures are treated in the same way, for several reasons. First, the existence of two distinct types of commands is confusing to the novice user. Second, *any* user becomes irritated when forced to type repetitive information, especially when the system already has it. Finally, the implementation of a given command may well change with time, typically from shell procedure to compiled program. This change could cause great pain to the users if it required the invocation method to change also.

3.2 Finding Commands

The shell normally searches for commands in a way that permits them to be found in three distinct locations in the file structure. It first attempts to use the command name without modification, then prepends the string "/bin/" to the name, and then "/usr/bin/". If the original command name is a simple one, the effect is to search in order the current directory, "/bin", and "/usr/bin". A more complex path name may be given, either to locate a file relative to the user's current directory, or to access one via an absolute path name.

This mechanism gives the user convenient execution of public commands and of commands in or "near" the current directory, as well as the ability to execute any accessible command, regardless of location in the file structure. The search order permits a standard command to be replaced by a user's command without affecting anyone else.

² As shown by a set of flag bits associated with the file.

3.3 Command Lines

A series of commands separated by "|" make up a *pipeline*. Each command is run as a separate process connected to its neighbor(s) by *pipes*, i.e., the output of each command (except the last one) becomes the input of the next command in line. A *filter* is a command that reads its input, transforms it in some way, then writes it as output. A pipeline normally consists of a series of filters. Although the processes in a pipeline are permitted to execute in parallel, they are synchronized to the extent that each program needs to read the output of its predecessor. Many commands operate on individual lines of text, reading a line, processing it, writing it, and looping back for more input. Some must read larger amounts of data before producing output; *sort* is an example of the extreme case that requires all input to be read before any output is produced.

The following is an example of a typical pipeline:

```
nroff -mm text | col | reform
```

Nroff is a text formatter whose output may contain reverse line motions; *col* converts them to a form that can be printed on a terminal lacking reverse motion, and *reform* is used here to speed printing by converting the (tab-less) output of *col* to one containing horizontal tab characters. The flag "-mm" indicates one of many possible formatting options, and "text" is the name of the file to be formatted.

A simple command in a pipeline may be replaced by a command line enclosed in parentheses "()"; in this case, another instance of the shell is spawned to execute the command line so enclosed. This action is helpful in combining the output of several sequentially executed commands into a stream to be processed by a pipeline. The following line prints two separate documents in a way similar to the previous example.

```
(nroff -mm text1; nroff -mm text2) | col | reform
```

If the last command in a pipeline is terminated by a semicolon (";") or new-line, the shell waits for the command to finish before continuing execution. It does not wait if the command is terminated by an ampersand ("&"); both sequential and asynchronous execution are thus allowed. An asynchronous pipeline continues execution until it terminates voluntarily, or until it is *killed* (by one of various means). A *command line* consists of zero or more pipelines separated by semicolons or ampersands.

For example, the following command line is used to run timing tests on an empty system. *Makeload* is a cyclic shell procedure used to generate a heavy, repeatable load of disk accesses, and *test1* performs timing tests on various programs. The shell runs *test1* with no load on the system, then starts one *makeload* to create a single unit of disk load for the second *test1*. Another *makeload* is invoked to yield two units of load for the last *test1*.

```
test1; makeload & test1; makeload & test1
```

Each *makeload* runs until explicitly killed by the user. A minimum of three processes are active by the time the final *test1* is run (two *makeloads* and one *test1*). In this particular case, all commands are implemented as shell procedures, so there is a separate invocation of the shell for each of the five commands on the above line, and each shell may well spawn hundreds of additional processes. Thus, a single user may consume all system resources by creating large numbers of long-lived asynchronous processes.³ More typical uses of "&" include off-line printing, background compilation, and generation of large jobs to be sent to remote computers.

3.4 Redirection of Input and Output

When a command begins execution, it usually expects three files to be already open, a "standard input," a "standard output," and a "diagnostic output." When the user's original shell is started, all three are opened to the user's terminal. A child process normally inherits these files from its parent. The shell permits them to be redirected elsewhere before control is passed to an invoked command.

An argument of the form "<file" or ">file" opens the specified file as standard input or output, respectively. An argument of the form ">>file" opens the standard output to the end of the file, thus providing a way to append data to it. In either output case, the shell creates the file if it did not already exist.

These forms of I/O redirection complement that of piping: files and programs can both be used as data "sources" and "sinks."

In general, most commands neither know nor care whether their input (output) is coming from (going to) a terminal, file, or pipe. Commands are thus easily used in many different contexts. A few commands vary their actions depending on the nature of their input or output, either for efficiency's sake, or to avoid useless actions (such as attempting random-access I/O on a terminal or pipe).

3.5 Generation of Argument Lists

Many command arguments are names of files. When certain characters are found in an argument, they cause replacement of that argument by a sorted list of zero or more file names obtained by pattern-matching on the contents of directories. Most characters match themselves. The "?" matches any single character; the "*" matches any string of any characters other than "/", including the null string. Enclosing a set of characters within square brackets "[...]" causes the construct to match any single one of the characters in that set. Inside brackets, a pair of characters separated by "-" includes in the set all characters lexically within the inclusive range of that pair.

For example, "*" matches all names in the current directory, "*temp*" matches any names containing "temp", "[a-f]*" matches all names beginning with "a" through "f", and "/u0/tnds/!f/bin/?" matches all single-character names found in "/u0/tnds/!f/bin".

This capability saves much typing, and more importantly, provides convenience in organizing files for various purposes. It allows convenient use of large numbers of small files.

3.6 Quoting Mechanisms

If a character has a special meaning to the shell, that meaning may be removed by preceding the character with a backslash ("\"); the "\" acts as an escape and disappears. Sequences of characters enclosed in double (") or single (') quotes are in general taken literally, except that substitution of shell arguments and variables is normally performed.

A "\" followed by a new-line is treated as a blank, permitting convenient continuation of multi-line commands.

4. USING THE SHELL AS A COMMAND: SHELL PROCEDURES

4.1 Invoking the Shell

The shell may be invoked explicitly in various ways:

sh - The new shell reads the standard input, but does not prompt. This is often used to let the shell act as filter, i.e., it can be used in a pipeline to read and execute a dynamically-generated stream of commands.

³ Lockout of other users in this way occurs several times per year on PWB systems; it is usually caused by overly enthusiastic beginning users.

sh file [args] A new instance of the shell is created to begin reading the file. Arguments can be manipulated as described in the next section

file [args] As noted in Section 3.1, if the file is marked executable, and is neither a directory nor a load module, the effect is the same as "sh file [args]".

4.2 Passing Arguments to the Shell

When a command line is scanned, any character sequence of the form \$n is replaced by the nth argument to the shell, counting the name of the file being read as \$0. A procedure may possess several different names and can check \$0 to determine the specific name being used, then vary its actions accordingly.

This notation permits up to 10 arguments to be referenced. Additional arguments can be processed using the *shift* command. It shifts arguments to the left; i.e., the value of \$1 is thrown away, \$2 replaces \$1, \$3 replaces \$2, etc. For example, consider the file "loopdump" below. *Echo* writes its arguments to the standard output; *if*, *exit*, and *goto* are discussed later, but perform fairly obvious functions.

```
: loop
if "$1" = "" exit
echo $1 $2 $3 $4 $5 $6 $7 $8 $9
shift
goto loop
```

If the file is invoked by "loopdump a b c" it would print:

```
a b c
b c
c
```

The form "*shift n*" has no effect on the arguments to the left of the *n*th argument; the *n*th argument is discarded, and higher-numbered ones shifted. Thus, *shift* is equivalent to "*shift 1*."

4.3 Shell Variables

Adding *if* and *goto* commands (described later) to the existing facilities permits convenient expression of some kinds of procedures: repetitive ones that perform a given set of actions for each argument and those that use simple conditional logic. Clear expression of many procedures requires at least a few shell variables.

The PWB shell provides 26 string variables, \$a through \$z. Those in the first half of the alphabet are guaranteed to be initialized to null strings at the beginning of execution and are never modified except by explicit user request. On the other hand, some variables in the range \$m through \$z have specific initial values, and may possibly be changed implicitly by the shell during execution. As will be seen later, few shell procedures ever use more than a few variables. A variable is given a value as follows:

```
= letter [argument]
```

If *argument* is given, its value is assigned to the variable given by *letter*. As an example of common usage, the procedure below expects to be called with a list of file names, optionally preceded by a flag "-w". If the first argument is "-w", the fact is recorded by setting \$a to "w", and the argument is shifted off the argument list, leaving only file names. If the first argument is not "-w", \$a is left unchanged, i.e., it is a null string.

```
if "$1" = -w then
    = a w
    shift
endif
```

... code to process file names, using \$a as needed

If no argument follows the letter in the "=" command, a single line is read from the standard input, and the resulting string (with the trailing new-line, if any, removed) becomes the value

of the variable. A common use is to capture the output of a program. For example, *date* writes the current time and date to its standard output. The following line saves this value in \$d.

```
date | = d
```

Thus, \$d would be set to a value such as:

```
Tue Jul 13 19:06:02 EDT 1976
```

A second use is in writing of *interactive* shell procedures, which are heavily used in PWB work. The following example is part of a procedure to ask the user what kind of terminal is being used, so that tabs can be set, delays changed, and other useful actions taken. The "</dev/tty" indicates a redirection of the standard input to the user's terminal; it is *not* seen as an argument to "=", but rather causes the variable to be set to the next line typed by the user.

```
: loop
echo "terminal?"
= a </dev/tty
if "$a" = ti goto ti
if "$a" = hp goto hp
    echo "$a no good; try ti or hp"
    goto loop
: hp
... processing for terminal type "hp"
exit
: ti
... processing for terminal type "ti"
```

Currently, five variables are assigned special meanings:

\$n records the number of arguments passed to the shell, not counting the name of the shell procedure itself. Thus, "sh file arg1 arg2 arg3" sets \$n to 3. *Shift* never changes the value of \$n.

\$p permits alteration of the names and order of directory path names used when searching for commands. It contains a sequence of directory names (separated by colons) that are to be used as search prefixes, in order from left to right. The current directory is indicated by a null string. The string is normally initialized to a value producing the effect described in Section 3.2: ":/bin:/usr/bin". A user could possess a personal directory of commands (e.g., "/u2/pw/bin") and cause it to be searched before the other three directories by using:

```
= p /u2/pw/bin : /bin : /usr/bin
```

\$r gives the value of the return code of the most recent command executed by the shell. When the shell terminates, it returns the current value of \$r as its own return code.

\$s is initialized to the name of the user's *login directory*, i.e., the directory that becomes the current directory upon completion of a login. Users can avoid embedding unnecessary absolute path names in their procedures by using this variable. This is of great benefit when path names are changed either to balance disk loads or to reflect project organizational changes.

\$t is initialized to the user's terminal identification, a single letter or digit.

In addition to these variables, the following is provided:

\$\$ contains a 5-digit number that is the unique process number of the current shell. In some circumstances, it is necessary to know the number of a process, in order to *kill* it, for example. However, its most common use to date has been that of generating unique names for temporary files.

4.4 Extended Order of Search for Commands

The user may request automatic initialization of each shell's $\$p$ by creating a file named `$.path` in the login directory. This file should contain a single line of the form shown for $\$p$. Every instance of the shell looks for this file and initializes its own $\$p$ from it, if it exists, otherwise `:/bin:/usr/bin` is used. Thus, the `$.path` value propagates through all of the user's shells, but changing $\$p$ in one shell has no effect on the $\$p$ of any other.

This facility is heavily used in large projects. It greatly simplifies the sharing of procedures, and can be quickly changed to adapt to changing organizational requirements.

4.5 Control Structures

The more complex shell control structures are actually implemented as independent commands that cooperate with the shell, but are not actually part of it. They are designed specifically for use in shell procedures, but are treated as ordinary commands. This separation of function allows the shell to remain a small program, efficient for on-line use, but still able to support powerful control structures in procedures.

4.5.1 Labels and Goto. The command `:` is recognized by the shell and treated as a comment. The most common use of `:` is to define a label to act as a target for *goto*. *Goto* is a separate command. Using `goto label` causes the following actions:

1. A *seek* is performed to move the read/write pointer to the beginning of the command file.
2. The file is scanned from the beginning, searching for `: name` on a line, either alone or followed by a blank or a tab.
3. The read/write pointer is adjusted (via the *seek*) to point at the line following the labeled line.

Thus, the only effect of *goto* is the adjustment of the shell's file read/write pointer to cause the shell to resume interpreting commands starting at the line following the labeled line.

4.5.2 If

`if expr command [args]`

If is also a separate command. If the conditional expression *expr* is found to be true, *if* executes the command (via *exec* system call), passing the arguments to it. If it is false, *if* merely exits.

The following primaries are used to construct the expression:

- `-r file` true if the file exists and is readable by the user.
- `-w file` true if the file exists and is writable by the user.
- `s1 = s2` true if strings *s1* and *s2* are equal.
- `s1 != s2` true if the strings are not equal.
- `n1 -eq n2` true if the integers *n1* and *n2* are algebraically equal. Other algebraic comparisons are indicated by `-ne`, `-gt`, `-ge`, `-lt`, and `-le`.

`{ command }` the command is executed; a return code of 0 is considered *true*, any other value is considered *false*. Most commands return 0 to indicate successful completion.

These primaries may be combined with the following operators:

- `!` unary negation operator.
- `-a` binary logical *and* operator.
- `-o` binary logical *or* operator: it has lower precedence than `-a`.
- `(expr)` parentheses for grouping. They must be escaped (as `\(` or `\)`), for example) to remove their significance to the shell.

All of the operators, flags, and values are *separate* arguments to *if*, and must be separated by blanks. The following are typical argument-testing operations

```

check a file argument to make sure it exists
if ! -r "$1" echo "can't read $1"

: assure either that: $1 is a and $3 is either b or c
: or that: $1 is d and $2 is e
if "$1" = a -a "(" "$3" = b -o "$3" = c ")" \
-o "$1" = d -a "$2" = e goto legal

```

Recall that the effect of the `\` at the end of the line is that of a blank. It is generally desirable to quote arguments when they might possibly contain blanks or other characters having special meaning to the shell.

4.5.3 If-then-else-endif. To improve the readability and speed of shell procedures, *if* was extended to provide the "if-then-else-endif" form. The general form is:

```

if expr then
... commands
else
... commands
endif

```

The "else" and the commands following it may be omitted, and it is legal to nest if's within the commands.

When *if* is called with a command, using the form of Section 4.5.2, it acts as described there, directly executing (or not) the supplied command. When called with *then* instead of a command, *if* simply exits on a true condition, allowing the shell to read (and interpret) the immediately following lines. On a false condition, *if* reads the file until it finds the next unmatched *else* or *endif*, thus skipping it and the lines in between. *Else* reads to the next unmatched *endif*, and *endif* is a null command.

These commands work together in a way that produces the appearance of a familiar control structure, although they do little but adjust read/write pointers.

4.5.4 Switch-breaksw-endsw. The *switch* command manipulates the input file in a way quite similar to *if*. It is modeled on the corresponding "switch" statement of the C language [RJT75A], and like it, provides an efficient multi-way branch:

```

switch value
: label1
... commands
: label2
... commands
.
.
: default
... commands
endsw

```

Switch reads the input until it finds:

- *value* used as a statement label, or
- "default" used as a statement label (optional), or
- the next unmatched *endsw* command.

Again, from the shell's viewpoint, the only effect of *switch* is to adjust the read/write pointer so that the shell effectively skips over part of the procedure, then continues executing commands following the chosen label or *endsw*.

Value is obtained from an argument or variable; if the label "default" is present, it must be the last label in the list; i.e., it indicates a default action to be taken if *value* matches none of the preceding labels. This construct may be nested; label enclosed by interior "switch-endsw" pairs are ignored.

The command *breaksw* reads the input until the next unmatched *endsw*, and commonly ends the sequence of commands associated with a label. *Endsw* is a null command like *endif*.

4.5.5 End-of-file and Exit. When the shell reaches the end-of-file, it terminates execution, returning to its parent the return code found in *\$r*. The *exit* command simply seeks to the end-of-file and returns, setting the return code to the value of its argument, if any. Thus, a procedure can be terminated normally by using "exit 0". The fact that *exit* is not part of the shell permits straightforward use of it as an argument for *if*.

4.5.6 The Missing Loop. Conspicuous by its absence is some form of *while* or *do*. All of the control structures described so far are implemented outside the shell; it appears that any useful looping construct requires significant changes to the shell itself. In any case, the most frequently observed kind of loop is that used to process arguments one at a time. For example, the following applies the first argument as a command to every remaining argument:

```

loop
if 0$2 = 0 exit
$1 $2
shift 2
goto loop

```

The "2" causes *shift* to leave the first argument in place.

4.5.7 Transfer to Another Command File—Next. The command "next name" causes the shell to abandon the current procedure and begin reading from file *name*. *Next* with no arguments causes the shell to read from the user's terminal. The idea of *next* is to permit the use of a file to initialize variables for use at the terminal:

```

= a /u2/pw/mash/articles
= b "nroff -rT2 -mm"
...
next

```

If this text were stored in file "init", it could be invoked by using "next init", causing the current shell to process it and return to the terminal. The user can then reference *\$a* and *\$b* appropriately. The user could of course use "=" to accomplish this directly, but at the cost of more typing.

Next is an attempt to obtain an effect like that of a subroutine call with a shared environment. It handles some problems well, but will probably be changed somewhat to make it more useful. Its most common application has actually been in very complex procedures that analyze their arguments, set up variables, then pass control to one of several successor procedures.

4.6 Interrupt Handling in Shell Procedures

Many PWB users have taken advantage of the ease and speed of writing shell procedures to automate various operations. In many cases, such procedures need to be used by clerical personnel who have no knowledge of these procedures' inner workings. A terminal interrupt (depression of "rubout" or "del" key) can be ignored or intercepted by a compiled program, or can cause termination of that program. The lack of interrupt-handling facilities in the shell quickly led to the usual problems:

- No procedure could use a terminal interrupt as a control mechanism
- Any procedure that created files for temporary use left them in existence if interrupted before it could remove them. In practice, any procedure that prints very much information is likely to be interrupted sooner or later.
- Some procedures need to temporarily ignore interrupts so they can guarantee consistency among files making up data bases. The PWB supports a profusion of packages that consist of file groupings accessed only through shell procedures.

The *onintr* command was added to solve these problems. It takes three forms: "onintr label" causes the effect of "goto label" to occur upon receipt of an interrupt; "onintr -" causes interrupts to be ignored completely; "onintr" alone causes normal interrupt action to be restored. A typical use of *onintr* is:

```

onintr cleanup
: create temporary file
ls -l | tee temp$$a | grep -c "^d" | = d
grep -c "-" temp$$a | = f
echo "directories: $d, files: $f"
: cleanup
rm temp$$a

```

This procedure displays the numbers of subdirectories and ordinary files in the current directory. The output of the *ls* command is a listing of the current directory; it is passed to *tee*, which makes an extra copy of it in "temp\$\$a", but also passes it to *grep*, which, in this instance, counts the number of lines whose first character is "d". This is the number of subdirectories, and is saved in variable *\$d*. The ordinary files (whose listing entries begin with "-") are counted in a similar way, and both counts are displayed. If the process is interrupted by the user, control transfers to "cleanup", where the temporary file is removed.

4.7 Additional Supporting Commands

4.7.1 Evaluation of Expressions. *Expr* supports arithmetic and logical operators on integers, and PL/I-like "substr", "length", and "index" operators for string manipulation. It evaluates a single expression and writes the result to the standard output, typically piped into "=" to be assigned to a variable. Typical examples are:

```

: increment $a
expr $a + 1 | = a

: strip off first 2 chars. of $1 and put result in $b
: expr substr abcde 3 999 returns cde
expr substr "$1" 3 999 | = b

: obtain length of $1
expr length "$1" | = c

```

The most common uses of *expr* are counting (for loops) and using "substr" to pick apart strings (such as the output from *date*, as in Section 4.3).

4.7.2 Echo. The *echo* command, invoked as "echo [args]", copies its arguments to the standard output, each separated from the preceding one by a blank, with a new-line appended to the last argument. It is often used to perform prompting or issue diagnostics in shell procedures, to add a few lines to an output stream in the middle of a pipeline, and to create editing scripts. "\n" yields a new-line and "\0n" yields the ASCII character given by the octal number *n*; "\c" is used to get rid of unwanted new-lines. For example, the following code prompts the user for some input and allows the user to type on the same line as the prompt:

```

echo "enter name:"
= a </dev/tty

```

4.8 Creation of Shell Procedures

A shell procedure can be created in two simple steps. The first step is that of building an ordinary text file, normally using the UNIX text editor *ed*. The second step is that of changing the *mode* of the file to make it *executable*, thus permitting it to be invoked by "name args", rather than "sh name args", as discussed in Section 3.1. The second step may be omitted for a procedure to be used once or twice and then discarded, but is recommended for longer-lived procedures.

The following shows the *entire* user input needed to set up a simple procedure to format text files according to a standard format and print the output on a particular type of terminal:

```
ed
a
nroff -rT1 -rC3 -mm $1 $2 $3 $4 $5 $6 $7 $8 $9 | gsi +12
.
w draft
q
chmod 755 draft
```

In the sequence above, the user called the text editor, entered a single line of text, wrote that line (creating the new file "draft"), and finally changed its *mode* to make it executable. The user may then invoke this command as "draft file1 file2", for example. The procedure calls the formatter with certain fixed arguments and any others supplied by the user; the formatter output is passed to *gsi* to convert it to a form that is appropriate for the user's terminal (in this case, a GSI-300).

If the sequence above were performed in a directory included in the user's ".path" file, the user could change directories and still use the "draft" command. Other people might make use of it also, especially if it were placed in a shared directory of commands.

The command sequence above could itself be stored as a shell procedure, although this particular sequence is an unlikely candidate for such an action. Note that the five lines following the *ed* call are processed by *ed* rather than the shell. It is quite reasonable to include data for other programs inside shell procedures, as long as those programs are careful in their method of reading, i.e., do not read beyond their own data. This method has beneficial results for performance, because I/O buffers can be shared without need for separate temporary files.

Shell procedures may be dynamically created by other shell procedures. A procedure may generate a file of commands, invoke another instance of the shell to execute that file, then remove it. An alternate approach is that of using *next* to make the current shell execute the new command file, allowing use of existing values of shell variables and avoiding the spawning of an additional process for another shell. In some cases, the use of a temporary file may be eliminated by using the shell in a pipeline. For example:

```
ls a* | sed "s/./cp & x&/" | sh -
```

The output of *ls* is a list of all file names in the current directory whose names begin with "a", one per line. *Sed* (a "stream editor") converts each line of the form "name" into the form "cp name xname",⁴ and passes it to a shell to be interpreted. A copy of each named file is generated under the name prefixed by "x".

Implied in the above discussion are several reasons why users like shell procedures better than compiled code. First, it is trivially easy to create and maintain a shell procedure, since it is only a file of ordinary text. Second, it has no corresponding object program that must be generated and maintained. Third, it is easy to create procedures "on the fly," use them, and remove them, without having to worry about managing libraries or about allocating disk storage. Finally, because such procedures are short in length, written at a high programming level, and kept in their source-language form, they are generally easy to find, understand, and modify.

⁴ *Cp* copies the file named by its first argument onto that named by the second, creating the latter, if necessary.

5. PATTERNS OF USAGE

5.1 Survey Methodology

A survey of PWB shell procedures was conducted, with the following goals:

- Discovery of procedures that could easily be redone in better ways, mainly to help users learn better ways of using existing tools, but also to improve system performance.
- Analysis of usage patterns, in order to improve existing PWB tools or build new ones. Rearrangements of command functions often occur when people recognize that the functions should be combined in different patterns. Examining user code is a good way to discover real needs.
- Determination of overall properties of procedures to help in redesign or enhancement to the shell and its supporting commands. This is part of an educational process aimed at putting existing tools to better use.

A program was written to find accessible shell procedures and print them for analysis by the author. Two facts made this a simple program. First, shell procedures are easy to manipulate because they are stored as simple text files. Second, the UNIX directory structure supports simple methods of traversal that permit easy investigation of *all* accessible files.

The files summarized in the next section represent a sample taken over a user population of more than 500 people. The data was collected in March, 1976.

5.2 Survey Results

A total of 1,725 procedures was analyzed. Table 1 summarizes the different forms of overall structure.

TABLE 1. Control Flow Summary

Category	Number	Percentage
1. single line	369	21
2. straight-line	935	54
3. argument loop only	83	5
4. branching, no loops	201	12
5. more complex	137	8

Procedures in category 1 consist of a single command line. Procedures in category 2 possess no control logic, but contain more than one line of text. The distribution of lines per procedure clearly favored small files, with a long "tail" of larger ones. This indicates that users find it helpful to "can" even one-line sequences. The large number of small procedures found on the PWB is at least partly due to the ease of creating them.

Thus, most procedures (75%) contain no control logic, but consist instead of straight-line scripts. The remaining 25% contain significant use of control logic. This percentage will probably increase with time, as more people become familiar with the shell's programming ability. In any case, it does indicate a significant need for control structures not provided conveniently in many CLS.

Procedures in category 3 are those whose control flow consists only of a single loop for processing the argument list, one at a time, such as:

```
: loop
if "$1" = "" exit
... commands
shift
goto loop
```

Procedures in category 4 are those possessing conditional branching, but no loops. Those in category 5 thus have at least one loop and at least one conditional branch in addition to the one implementing the loop.

Each entry in Table 2 gives the number of shell procedures and the percentage of the total in which the specified construct was found. These figures were obtained by visual inspection of the files, and should probably be taken as lower bounds on the number of occurrences.

TABLE 2. Occurrence of Shell Constructs

Shell Construct	Number	Percentage
switch (or obvious use for it)	45	3
non-shell commands in file	514	30
pipe into shell (command sh -)	32	2
parenthesized commands (i.e., implicit sh call)	39	2
explicit sh calls or obvious shell procedure calls	170	14
= a </dev/tty (read line from terminal)	81	5
command = a (assign output of command to variable)	124	7
= a string (assign string to variable)	146	8
onintr label (intercept interrupt)	35	2
onintr - (ignore interrupt)	14	1
expr substr	33	2
expr + or -	24	1
next	13	1

The figures indicate the frequent intermingling (30%) of shell and non-shell commands in the same file, utilizing the fact that the shell and other programs share the same input. This intermingling occurs most often in the straight-line scripts. At least 14% of the procedures invoke another level of shell, in one or more of the three ways listed. When CL subroutines are available, people do make use of them, and tend to become extremely irritated with any CL that does not provide them. At least 5% of the procedures expect to write prompts to the user and read from the terminal; the figure would be much higher if it included prompts issued by programs other than the shell.

Despite the fact that *expr* supports multiplication and division, no serious uses of these operations were found.⁵ From observation of the procedures, it is clear that users are doing much more string manipulation than arithmetic. Although *if* can perform arithmetic comparisons, the bulk of the operations performed by *if* were string comparisons. Likewise, the string operations of *expr* occurred more frequently than the arithmetic ones. It appears that CL arithmetic is a helpful facility, but string manipulation is a necessity, an ordering contrary to that of some existing CLS (COL76A).

5.3 Informal Observations

In many cases, a shell procedure is kept in a directory with the files that it manipulates, and is never used except in that directory. This is convenient in practice, and allows the user to forget about possible naming conflicts with procedures in other directories. It also supports the common practice of packaging related files in a separate directory.

⁵ Multiplication was used in two procedures, but they were ignored in the survey because they were obviously "play" procedures, i.e., two versions of a procedure to compute factorials, one iterative and the other recursive.

Multi-person projects often possess several sets of directories used to store shared procedures, and individuals may well have their own directories in addition. For example, project "tnds" may use "/u0/tnds/bin" as a repository for procedures needed by the entire project. Group "tf" within "tnds" may use "/u0/tnds/tf/bin" for procedures needed only by the group, and individual "jtb" might use another directory for personal commands. The ".path" file used in this case might contain:

```
:/u0/tnds/tf/jtb/bin:/u0/tnds/tf/bin:/u0/tnds/bin:/bin:/usr/bin
```

In general, people use this type of sharing mechanism to adapt the system to their own organizational needs.

5.4 Command Usage

Although the survey concentrated on static analysis, a few dynamic usage statistics may be of interest. Each of the larger PWB computers execute 25,000-40,000 commands per day. Approximately 80% of these commands are executed within shell procedures. These figures provide clear evidence that people utilize the shell programming capabilities to a very large degree.

6. AN ASSESSMENT

In the environment described in this paper, advances occur in three distinct ways: improving old tools, building new ones, and improving the framework used to put them together. Work is under way in all three areas, including most of the following:

- Addition of *while* to the shell.
- Creation of a "high-speed" shell that includes the bulk of the control structure code.
- Addition of different methods of manipulating pipes. For example, one would like to use programs that can have several input or output pipes, rather than only one of each.
- Cleanup of the overall syntax, which has its ugly aspects.
- Better ways of handling variables, including more powerful default and substitution mechanisms.
- Commands to perform argument processing without loops.

It is clearly recognized that the PWB CL does not do everything that its users would like it to do, so it is likely to continue evolving at a rapid rate. This paper has attempted to present a consistent view of the PWB CL as it has existed for a relatively long time (about nine months). It has emphasized the description of existing facilities whose usage patterns are well known, rather than the forecasting of things to come.

7. CONCLUSIONS

This paper has described a minicomputer-based CL that is more powerful and convenient to use than many of those available on much larger systems. It is a real CL used daily by people doing production work under tight schedules. Although much work remains to be done on it, it has definitely shown its worth as a programming language for many applications, and has made major contributions to the ability of users to get their work done with a minimum of effort and irritation.

ACKNOWLEDGEMENTS

The shell was originally written by K. Thompson; that it has been so adaptable to a wide variety of uses and machine configurations is a great tribute to its creator. Many of the additions were implemented by R. C. Haight and A. L. Glasser.

REFERENCES

All references cited in this paper appear at the end of "An Introduction to the Programmer's Workbench," by Dolotta, T. A., and Mashey, J. R., in these Proceedings.

Documentation Tools and Techniques

J. R. Mashey
D. W. Smith

Bell Laboratories
Piscataway, New Jersey 08854

Keywords: Text processing, text formatting, UNIX.

Abstract: In a software development project of any appreciable size, the production of usable, accurate documentation may well consume more effort than the production of the software itself. Several years of experience on many Programmer's Workbench projects have shown that document preparation should not be separated from software development and that the combination of a flexible operating system, a powerful command language, and good text processing facilities permits quick and convenient production of many kinds of documentation which might be otherwise unobtainable, impractical, or very expensive. Our basic approach has been to develop techniques for effective combination of existing UNIX facilities. A number of case histories are given to illustrate the flexibility, convenience, and general usefulness of these techniques.

1. INTRODUCTION

The Programmer's Workbench (PWB) [DOL76A] is a working example of a system specialized for program development. Many activities performed with the PWB's help can be considered to be in some sense document preparation or text processing:

- editing source code or documents;
- scanning output returned from other systems;
- scanning files for certain data (e.g., cross-referencing, data collection and summarization);
- recording changes to source or document text using SCCS [ROO75A];
- preparing scripts for load or regression tests [DOL 5B];
- examining output from load or regression tests;
- preparing input for and producing many kinds of reports;
- entering and maintaining information on activities and schedules.

From our observation of actual PWB usage, it appears that flexible, integrated text processing is an absolute requirement for any program development facility. This requirement is *not* satisfied by adding a simple text formatter to an existing system; the *entire* system must support text manipulation in a consistent, convenient way. Text formatting becomes more flexible and adaptable when implemented as part of a more general environment, rather than as an isolated, specialized "add-on." In addition, many normal programming activities are better supported when they have good text processing facilities available in the same way as other tools. Specific advantages include:

- The same editor is used to enter and modify documentation, source programs, and procedures for manipulating them both. This unification helps lower training and support costs.
- Likewise, it is possible to use existing tools to manipulate either source programs or text without regard for the intended use of these tools. Obvious examples include file scanners, sorts, and cross-reference commands. The PWB user has access to all the tools provided by UNIX [RIT74A]. The set of tools is extensive [KER75A], and unexpected uses are constantly found for them.
- When a document must include samples of source programs, it is easy to assure accuracy of such samples because they can be automatically copied from actual source program files.

- No artificial distinctions need exist among source programs, documentation, and data bases. A set of files and command procedures may appear to be a small interactive data base management system; viewed another way, it is just a mechanism for producing a large, complex document. Finally, information from these files may be used to automatically generate or verify parts of source programs.
- It is often *more* important to automate the management and control of text than it is to format it. It is especially necessary to be able to do this quickly and cheaply, because procedures and needs not only differ among projects but change quickly and often.

In this paper we illustrate some of the facilities one gains by utilizing a concept unique to the UNIX system—that text processing is an integral part (not an "add-on") of an effective program development environment. We first briefly survey some of the useful tools provided by UNIX. We then show the desirability, convenience, and (in our opinion) the *necessity* of this approach by examining case histories of several projects implemented on the PWB in the last few years.

2. BACKGROUND

2.1 Environment

Many aspects of the UNIX environment make it especially suitable for text processing.

First, a UNIX file is just a string of characters, whose format and interpretation are under control of the user, rather than UNIX. The most common file format is that of an *ASCII file*, i.e., a sequence of arbitrary-length *lines* terminated by *new-line* characters. This format avoids any need to worry about blocking or blank stripping. Explicit allocation or de-allocation of storage is avoided by the ability of files to grow and shrink dynamically.

Second, UNIX supplies a flexible hierarchical file structure that is easily adaptable to user needs. Convenience of document organization is improved by the ease of grouping large numbers of files. Users may share files and commands in various ways.

Third, the UNIX *pipe* mechanism is particularly desirable for text processing because it provides an effective way of manipulating data streams. A pipe connects two executing programs, causing the output of the first to be read as the input to the second, while all synchronization and buffering details are handled by UNIX itself. Two or more programs can thus be connected together to form a *pipeline*, an object commonly found in UNIX text processing applications.

Finally, although some types of terminals are more convenient than others for UNIX usage, UNIX is easily adaptable to many different kinds of terminals. It is rarely necessary to modify the operating system to take advantage of some new type of terminal. All that is usually necessary to take advantage of any unusual terminal features is to write a single, small, non-privileged command to cater to the specific terminal.

2.2 Specific Tools

Although almost all UNIX commands turn out to have some text processing use sooner or later, some are especially useful.

2.2.1 Editing System. *Ed* is the standard editor. A related command, *ved*, is a *stream editor*, whose use is more efficient and convenient than that of *ed* in one pass applications, i.e., when used in a pipeline.

2.2.2 Text Formatting Software. The simplest version of the "run off" formatting software is *roff*, whose formatting power is roughly equivalent to that of IBM's AIS, although its input appearance is radically different. *Nroff* is much more powerful than *roff*, permitting the use of sequences of formatting codes (known as macros) whose behavior may be modified by arguments specified when such a macro is interpreted. The typesetting software, *troff*, is similar to *nroff* with respect to programming, but its output is intended for a Graphic Systems, Inc., phototypesetter, which can produce various font styles and character sizes as well as numerous special characters for mathematical expressions and for other purposes.

2.2.3 Output Filters. Most *roff* and *nroff* output can be printed on most terminals, but some output takes advantage of special features and escape sequences found only on the Teletype® Model 37. To make certain sequences (such as reverse line feed, half-line motion, or extended character sets) function properly on other output devices, the output of *roff* or *nroff* must be processed ("filtered") by other programs before being physically printed. For example, *gsi* is used to make output acceptable to certain terminals that use the Diablo HyType® print mechanism; *hp* is used for output on a Hewlett Packard 2640A. *Col* is used to "normalize" such output for terminals lacking reverse line feed, especially when multi-column output is desired.

2.2.4 Input Filters. Just as output filters process output before it is printed, input filters can process text before it is formatted. Such filters often handle complex or unusual requirements without requiring modification of the formatters. For example, *eqn* and *neqn* permit convenient formatting of mathematical expressions [KR75B], and *tbl* helps align and format tabular matter and its associated headings.

2.2.5 Data Storage Facilities. An *archive file* is a collection of files created and maintained via the *ar* command. It is often used to store large numbers of small files in an efficient way, and may be used in some data base applications. The commands that are part of SCCS [ROC75A] can be used to maintain control over document changes, record such changes, and recreate the document as it existed at any point in time.

2.2.6 File Examination Software. The *grep* program is used to extract from text files lines containing a specified sequence of characters. *Cref* produces cross-reference listings of the C and UNIX assembler programming languages. The *diff* program reports differences between two files, listing the lines that have been added, deleted, or changed. The *comm* program lists lines that are common to two files. *Typo* helps the user quickly identify possible typographical errors.

2.2.7 Command Language. The UNIX command language is interpreted by a program called the *shell*. It provides a quick, inexpensive way to combine other tools into the desired procedures [MAS76A]. For example, the procedure in Figure 1 formats a set of files according to a standard *nroff* macro set "-mm", allowing an optional first argument "-g" to request use of the *gsi* filter and 12-pitch output; otherwise output is produced for a terminal without reverse-line motion (e.g., a G.E. TermiNet 300).

```
if $1 = -g then
  shift
  nroff -h -rT1 -mm $1 $2 $3 $4 $5 $6 $7 $8 $9 | gsi +12
  exit
endif
nroff -mm $1 $2 $3 $4 $5 $6 $7 $8 $9 | col
```

Figure 1

3. PROGRAMMER'S WORKBENCH MEMORANDUM MACROS

The *nroff* and *troff* are unusually flexible and powerful text formatters. However, to make the best use of this software, one must become quite familiar with all the basic formatting requests. Learning the requests and the useful ways of combining them is a full-time *programming* task. For the PWB, we spent considerable effort in examining what kinds of things people would want to do in manipulating and stylizing their documentation—whether preparing letters, technical reports, user manuals, or software documentation. With the design priorities as described below, we developed the Programmer's Workbench Memorandum Macros (PWB/MM), a unified, consistent, and flexible set of formatting codes. PWB/MM minimizes the differences between *nroff* and *troff*, making it easier to enter a document without concern for its eventual output medium. This package brings much of the inherent power of the formatters to the user who does not want to spend a great deal of time learning their intricacies. It is heavily used by clerks and typists, as well as by technical personnel. PWB/MM was used to prepare this paper, as well as the five companion PWB papers [BIA76A, DOL76A, DOL76B, KNU76A, MAS76A].

3.1 Design Priorities

The following are the qualities we emphasized in the design and implementation of PWB/MM, in approximate order of importance:

- *Robustness in the face of error*—When the input is incorrect, either a reasonable interpretation is made of an error, or a message is produced when an error is detected.
- *Ease of use for simple documents*—It is not necessary to specify complicated sequences to produce simple documents, especially those following any reasonable format. Default values are provided where at all possible.
- *Parameterization*—People have many different preferences in the area of document styling. Many parameters are provided so that users can customize output to their needs over a wide range of styles. These parameters are set once at the beginning of the document. The formatting codes are entered consistently, regardless of the parameter settings.
- *Extension by moderately expert users*—We have made a strong effort to use mnemonic naming conventions and consistent techniques in the construction of the macros. A user can add new macros or redefine existing ones if necessary. There is also the provision, in certain cases, to have user-defined macros called at appropriate points within the provided macros.
- *Device independence*—The most common use of PWB/MM is printing documents on hard-copy typewriter-like terminals, using *nroff*. Output can be printed on either 10-pitch or 12-pitch terminals and can be examined on an appropriate CRT terminal. Finally, phototypeset output can be produced using *troff*.
- *Minimization of input*—The design of the macros lessens repetitive, unnecessary typing. For example, to obtain a blank line after all first or second level headings, the user need only set a specific parameter once at the beginning of a document, rather than add a blank line after each such heading. A table of contents, as well as other information derived from the input text, can be generated automatically.

3.2 Features

PWB/MM features permit the user to concentrate on the *logical structure* of the document, not on its eventual *appearance*. We feel this is a desirable direction of evolution for text processing. Some specific examples include the implementations of headings, various styles of lists, and footnotes.

3.2.1 Heading Styles. The formatting code ".H" specifies a new heading. The arguments to this code indicate the level of the heading (first or top level, second level, etc.) and the text of the heading itself. For example, the two headings immediately above were generated by:

H 2 Features

H 3 "Heading Styles."

One does not need to specify the number to be assigned to a heading. Headings are automatically numbered when the document is formatted. Thus, if sections of text are rearranged (which is very easy to do with the UNIX editor *ed*), one need not worry about renumbering any affected headings.

The default heading style produces numbered headings, as illustrated by this paper. By setting some parameters before the first heading, the user can easily modify this style to obtain, for example:

- any of various outline styles, replacing the default sequence "1., 1.1, 1.1.1" by, for instance, the sequence "I., A., 1.;"
- centered headings in place of left-justified ones;
- different fonts (or underlining style) for each level;
- different choices of pre- and post-spacing.

What must be noted is that although the style may vary, the way of typing a heading does *not*. A few global parameters control the final overall appearance.

This approach not only contributes to uniformity of style within a document, but also allows the user to make radical changes in style *after* the document has been entered. Finally, the same text can be included in several documents that must adhere to differing standards, as in the case when an internal report is submitted to a journal that requires another format.

3.2.2 Lists Various kinds of lists are often needed:

Ordered lists	Items within the list ("list items") are numbered or "lettered" sequentially.
Unordered lists	Items in the list are marked in some way, say, with a dash or a bullet.
Definition lists	Each item in the list is a word or a phrase followed by an explanation or definition. (This list is an example of a definition list.)
Reference lists	Lists of references used in the body of the text.

As an example of the ease with which such lists can be entered, Figure 2 shows the input text for the above list.

```
.VL 18
.LI Ordered\ lists
Items within the list ("list items") are
numbered or "lettered" sequentially.
.LI Unordered\ lists
Items in the list are marked in some way, say,
with a dash or a bullet
.LI Definition\ lists
Each item in the list is a word or a phrase
followed by an explanation or definition.
(This list is an example of a definition list.)
.LI Reference\ lists
Lists of references used in the body of the text.
.LE
```

Figure 2

All formatting codes begin with a period in column 1. Each kind of list in PWB/MM consists of the following three parts—a "list begin" which specifies the type of list (.VL in this case), one or more "list items" (.LI followed by the text), and the "list end" (.LE) marking the end of the list. All information about the kind of list appears in the "list begin." It is very simple to add, delete, or reorder the items within a list. For ordered lists, the numbering or lettering is handled automatically so that, in inserting or deleting items, one need not manually renumber any of the items. Lists can be nested to a depth of six levels.

3.2.3 *Footnotes.* To enter the text of the footnote and to generate the next footnote number, one need only type *F after the word to be noted, followed by the text of the footnote delimited by two formatting commands. For example, the following input text:

```
This is the word\F
.FS
The text of the footnote
goes here.
.FE
that is to have a footnote.
```

will produce the following output:

This is the word¹ that is to have a footnote.

as well as the footnote at the bottom of this page.

PWB/MM automatically records the footnote for subsequent inclusion at the bottom of the current page. If there is not enough room on the current page for the entire footnote, the remainder of the footnote is saved and printed at the bottom of the next page. All the mental processes of the typist in handling footnotes (gauging the length of the footnote text, remembering to bring up the bottom margin by the proper amount, etc.) are handled automatically.

3.3 PWB/MM Summary

PWB/MM supplies many services in addition to those described above. The end user obtains the benefits of complex formatting procedures without having to learn the techniques of their implementation. Note that PWB/MM represents a *layered* approach to text formatting: it is a second layer built on *nroff* and *troff*. We feel that a common formatter design error is that of building *too much* into the formatter itself, rather than using a layered approach to formatting complex documents. For example, the fact that the footnote mechanism was not locked into our formatters allowed us much more flexibility in creating variant footnote styles.

4. CASE HISTORIES

The UNIX environment permits quick construction of working systems from existing components, easy adaptation of existing files to new purposes, and inexpensive modification of systems to meet changing needs. To support these assertions, we describe very briefly a few applications chosen from the many that have been built in the last two years.

4.1 A Text Processing Application

One of the Bell System companies had a special text processing requirement. The company conducts a weekly management seminar and is required to prepare timely reports describing each session—what material was presented, what comments were made, a list of any unresolved questions, etc.

4.1.1 *Text Production and Analysis.* The initial need of the seminar is a text processing facility—the ability to enter, update, and print various kinds of reports for each session. The reports are printed on hard-copy terminals for editing and author review. The final reports are phototypeset for distribution to top management.

4.1.1.1 *Report Production.* The reports are entered using the UNIX editing system, embedding within the text of the reports certain formatting codes. A precursor to PWB/MM, these formatting codes were designed to be used for either *nroff* or *troff*. The *typo* program is also utilized to help find typographical errors.

Each session closes on a Friday; by the following Wednesday, a 70-page booklet containing the phototypeset originals of all the material from that session is ready for reproduction. The staff in

¹ The text of the footnote goes here

charge of the seminar has been quite pleased with the documents produced. They have stated that there would have been no other way to produce the same high-quality output in the same period of time.

4.1.2 Analysis of Content. With several weeks' data stored on the system, the question arose as to possible ways of extracting portions of that data. It appeared reasonable to use the system to produce another report consisting of all paragraphs containing the word *marketing*, for example

This analysis was accomplished by considering the structure of the text in the report files. Due to the use of the formatting codes, it turned out that all sections of text (paragraphs, items in lists, etc.) were preceded and terminated by specific character strings. Although a paragraph might contain the context word (like *marketing* or *Marketing*) several times, such a paragraph is to appear only once in the output report. Thus, it was necessary to use the editor and the *sort* command to obtain a unique list of the line numbers of "paragraph beginnings." Another editor script alters this list of unique numbers to obtain yet a third editor script that causes all the lines of the paragraph to be printed, taking care of cases such as paragraphs that are followed by lists. Another pass through the original file using this last *generated* script produces the desired output.

Once these scripts were developed, it became a very simple matter to put them together in a *shell* procedure. All that the user needs to specify in invoking this procedure is the name of the file containing the data and the context word by which the paragraphs are to be selected.

4.1.2 Conferee Roster. Another text processing task is to maintain a list of conferees scheduled to attend the seminar. This list is periodically updated since persons are often rescheduled due to business demands. The roster of conferees for a given session might not actually be finalized until a few days before the beginning of the session.

The data for all the conferees attending a given session are entered into the same file, whose name is the week during which the session is held. The data for each conferee is entered as shown in Figure 3, using formatting macros for each item of information needed. The "company" data includes special two-letter codes for all Bell System companies.

```
last name      .LN Smith
first name     .FN "Dale W."
title         .TL "Member of Technical Staff"
company        .CO \*(BL
address        .AD "6 Corporate Pl., Piscataway, NJ"
telephone     .TN "201 981 7315"
department    .DP "Support Products and Systems"
level         .LV 0
name/address flag .NA
```

Figure 3

After several sessions of the seminar, the accumulated rosters took on added importance. It is possible to scan the data to determine whether or not a particular person has ever attended the seminar and, if so, the week of attendance. For instance:

```
grep "\.LN.*name" filenames
```

will list all files containing an entry for a person whose last name is *name*. Providing this information by manually scanning printed rosters is time-consuming and error-prone.

With this data, reports listing such information as:

- Who has attended from the marketing division?
- How many fifth level managers from XYZ Company have attended?
- What engineering personnel from XYZ Company have attended?

are very easy to obtain. What is needed is a way to say "Give me all the data about conferees who have this property." This is accomplished by having the user specify, to a shell procedure, what the desired property is (utilizing, of course, information about the structure of the data). To obtain all the data about the conferees from any given week who are in the marketing department, one would execute

```
ed - filename -output
g/\DP.*Marketing/?\LN?/\NA/p
q
```

By first placing all the conferee data into a temporary file and then executing the above editor script on that file, one obtains the data about *all* conferees from marketing.

One can further categorize the extracted data by executing a similar editor script looking for the desired property. For example, if the file *mkt* contains all the marketing conferees, one could obtain all fifth level marketing people by

```
ed - mkt >fifth
g/\LV.*5/?\LN?/\NA/p
q
```

By providing alternate definitions for the formatting macros, one can generate a data record for each conferee. These records can then be sorted in various ways (by company, by name, etc.). Again, the ability to easily combine existing UNIX facilities proved invaluable in extracting additional information from existing data.

4.2 A Data Base Application

Any programming project that involves more than a few people seems to generate innumerable small data bases to record information needed for documentation and control of the project. Such data bases are often small (a megabyte or less), need to be implemented in a timely fashion, and vary greatly in nature from project to project. To handle this sort of data base, one can either utilize an existing general-purpose data base management system (DBMS) or write several programs to implement the desired application. The former approach often yields a severe form of overkill, while the latter may require too much programming effort and unacceptable delays.

A third approach is that of building an extremely specialized DBMS for the specific application. Programming costs and lead time are minimized by writing command language procedures that combine existing tools in appropriate ways. The advantages of this approach are: speed of implementation, satisfaction of specific needs of the application, small size of code (no compiled code at all), ease of maintenance, reliability, and good output appearance (e.g., by using *nroff*). The main disadvantage is the possible lack of execution speed. Also, there may be some inconvenience in handling larger data bases or those whose structure has a bad match to the UNIX file system. On the whole, this approach appears to be quite worthwhile; it has been applied very successfully a number of times.

As an example, the first such system known to the authors was built for a large programming project during September, 1974. The project had been using an obsolete DBMS to store information regarding error messages produced by the project, and a better DBMS was desperately needed. The project includes a group of 400 modules, each of which can produce error messages of two different types. Certain information must be maintained for each error message, such as the cause of error, actual message text, response required, responsible person, etc. There are about 3000 messages altogether, requiring about a megabyte of disk storage, organized into 550 UNIX files. One author (Mashey) spent about three work-days (spaced over a period of two weeks) as follows:

- Discussions were held with project members to determine their requirements.

- Data formats were designed
- Procedures were written to add, update, delete, and print individual items in the data base
- Procedures were built to produce various kinds of summaries and large reports
- Procedures were included to maintain logs of data base changes
- All of the procedures were tested and (more or less) debugged

By using existing tools and by modifying the package according to changing user perceptions, a usable system existed at the end of two weeks. Over the next few months, the data base was loaded with information from the old system—the major part of the time being required for data purification. The total time spent by the author on this project was about two work-weeks, of which half was documentation time. By February 1975, the entire project had been turned over to a relatively new UNIX user who has had little trouble in modifying and extending it. At turnover time, the entire package of procedures and formatting macros consisted of nothing but text files, occupying a total of 30,000 bytes.

This system is used in various ways. In one sense, it is a small DBMS; in another, it is just a way of manipulating a large, complicated document. Finally, it is also used to verify the consistency of the actual programs with their documentation. A selected subset of the information is extracted and sent to the machine on which the application project executes. It is then compared with the actual message text.

It is interesting to note that data base management software *per se* has not proved popular on UNIX to date, because more flexibility can be obtained through the appropriate use of *sort*, *ed*, and other UNIX commands.

4.3 Production of Revision Bars

Existing tools can often be combined to produce novel results, although sometimes a new tool must be built to fill a gap between some others. An example is the process used to generate automatic revision bars (or other marks) for modified documents. When manuals are revised and distributed to users, such revision bars are placed in the margins to indicate the locations of changed text. The formatters support a mode of operation in which a specified character is placed in the margin of every line of output text. This mode may be turned on and off by requests in the formatter input. The only problem is that of determining where to insert these requests. The following command line illustrates the entire process:

```
diff - old.new | diffmark temp | ed - old; nroff temp
```

Diff compares the two files, and generates an editor script to convert file *old* into file *new*. The resulting script is dynamically piped ("|") into *diffmark*, which rearranges this script and then passes it to the editor. The editor, in turn, edits the *old* file, but writes the result into *temp*. *Temp* is identical to *new*, except that formatter change-mark requests are included in appropriate places. *Nroff* then formats the result with the needed revision bars.

Diffmark is a small, simple program written specifically to bridge the gap between the other commands and was intended

only for use with documentation. As commonly happens, a number of people immediately applied the same process to other kinds of text, i.e., source programs and data base schemas.

In some cases, a user simply makes a copy of a file before editing it, performs the process described, and removes the *old* and *temp* files after *temp* has been printed and distributed. If more extensive control is desired, SCCS [ROCKSA] can be used to maintain complete control of change history and recreate any needed version of a file. This is a good example of the unexpected advantages of keeping documentation and source together in an environment where they can both be handled with the same tools.

4.4 Miscellaneous Applications

4.4.1 Viewgraph Production. With the phototypesetting software *troff* and a set of formatting macros developed for the PWB, one can easily prepare viewgraphs for lectures, demonstrations, etc. The advantages for doing this are as previously stated—the ability to easily modify the data, the ability to include actual programming statements, and the assurance of accuracy during revision.

4.4.2 Bibliography Management. To avoid wasted effort when papers are written, especially by groups of people, a bibliography data base system was built from ordinary text files and a few shell procedures. A unique key is assigned to each reference as it is added to the data base. A document is scanned for keys of this form, and a list of unique keys extracted and sorted. The list is used to search the bibliography data base and extract the items to be included in the reference list. The procedure to perform all of these actions was written, coded, and debugged in 20 minutes. It was used to simplify the preparation of the references for this and the related papers.

4.4.3 Milestones and Scheduling. A need arose for a simple, easy to use way of managing schedule dates and activities for individuals and groups of various sizes. It was an afternoon's work to create a package of formatter macros, text files, and shell procedures to accomplish this. Schedules can be summarized at various levels, and ordered according to half a dozen sort keys. The entire package consists of about six pages of text.

5. CONCLUSIONS

We have stated the opinion that a software development facility needs a good, integrated set of text processing tools. We have summarized the tools most commonly used, and described a few of their applications to support our opinion. They are but a very small sample of those available; UNIX makes it almost too easy to invent and implement useful applications on the spur of the moment. From our experience, this a very desirable situation since many useful, effective aids have been quickly and cheaply built this way. In many other environments, they might well be too small, too difficult, or too expensive to ever be done.

REFERENCES

All references cited in this paper appear at the end of "An Introduction to the Programmer's Workbench," by Dolotta, T. A., and Mashey, J. R., in these Proceedings.

A User's Viewpoint on the Programmer's Workbench

M. H. Bianchi
J. L. Wood

Bell Laboratories
Piscataway, New Jersey 08854

Keywords: Software development, programming aids, UNIX.

Abstract: The Programmer's Workbench boasts a broad set of highly useful features aimed at the application program developer. It claims to be a "human-end" computer providing tools and services to ease the load on the application system designer, programmer, documenter, tester, and delivery personnel. This paper shows the benefits of using the PWB tools, individually and in combination. Through specific examples drawn from the history of a software project, evidence is given that the use of the Programmer's Workbench can be a major contributing factor in the successful development of a software project.

1. RELATION OF CDS DEVELOPMENT GROUP TO PWB

The Circuit Design System (CDS) group was getting ready to write code at about the time that the Programmer's Workbench (PWB) was starting to accept customers. An informal arrangement was made to allow the authors to try out the new system. After about a month of discovering the tools that PWB offered, the arrangement was made official. For a while, CDS was the heaviest application development user on the Workbench, and hence we were the first to ask many questions and make many comments.

We have seen the Workbench grow and have been users for over two years. Unlike the developers of the PWB, we are only users. We will demonstrate through discussion of our experiences that the Programmer's Workbench *concept* is viable. Moreover, we will show that the actual PWB at Bell Laboratories is a most important contributor to the successful development of a project that ultimately runs in an unrelated environment.

Many of the PWB's facilities can be found on other systems in some form or other. From the user's viewpoint, the PWB provides an unusual variety of program development tools in a single, uniform, and easy to manage environment. This paper is *not* intended as a catalog of new or exotic facilities, but as a summary of one group's experience in using the tools provided.

We will be talking about PWB strictly from the viewpoint of a user who does not see, and is generally unconcerned with, the details of PWB implementation. The reader should be familiar with [IXL76A], which provides an overview of the PWB and a rationale for its existence. [RIT74A] describes UNIX, the time-sharing system on which the PWB is based. Finally, the discussions in [MAS76A] may improve the reader's understanding of some of the more complex examples presented here.

2. CDS—A QUICK OVERVIEW

The Circuit Design System mechanizes certain functions performed in the day-to-day activities of a Bell System operating telephone company. It uses other software systems written at Bell Laboratories to provide data base information, but its own emphasis is strongly in the engineering field. During the development cycle, the majority of our personnel were communications engineers and not data processing professionals.

The system itself must coexist with another system that utilizes IBM's Information Management System (IMS) to provide hierarchical data base management and transactional telecommunications [IBM75B].

The purpose of the first development cycle was to test the feasibility of the engineering process. Therefore, we had two secondary objectives. The first was to use the cheapest equipment possible, and the second was to minimize overall experimental costs. Our final product contained 196 PL/I program modules and 6 data bases accessed from dial-up terminals.

The developers themselves had to perform all the tasks involved in program maintenance. PWB allowed us to set up procedures that drastically reduced the amount of time required to maintain the system. In many instances, whole tasks, such as partitioned dataset compression, were made totally automatic.

3. ENVIRONMENT BEFORE PWB

The primary program development tools available at our location were Applied Data Research's LIBRARIAN [ADR73A] and IBM's standard utilities [IBM72A]. The source editing features of these programs do not lend themselves to making complex updates to modules. Moving blocks of code from one section of code to another is almost prohibitively difficult. Thus, as a program is modified, it becomes riddled with branches that have nothing to do with the implementation of the algorithm. Also, it is a batch system and an error in editing can ruin a half-day's work.

4. EARLY PWB TOOLS

The first version of PWB to which CDS was exposed was a DEC PDP-11/45 running UNIX plus a Remote Job Entry (RJE) capability. Many of the programs that were to enhance the PWB concept were still in development.

But the early support provided by this one system was of great value to us. We found it relatively easy to use, extremely reliable, and adaptable to many of our needs with little effort.

4.1 The Text Editor

At first, the major tool used was the UNIX text editor, *ed*, with its very terse syntax and surprising flexibility. Previous experience with the QED style of editor was a definite advantage in learning about *ed*. The first real work done with *ed* was to enter two small PL/I programs for use in the CDS project. These were thought out and entered at the terminal by the programmer. This early exercise convinced us that *ed* was going to be a valuable tool. We were impressed by the ease of editing and moving code around, the time saved by entering code directly into a computer rather than using coding forms and keypunches, and the ability to "desk check" while at the terminal.

4.2 The UNIX File System

UNIX presented us with a true tree-structured file system that allowed us to build logical relationships between its files and directories (leaves and nodes).

The CDS directory became the root of our "program tree" which we present in part in Figure 1. We built personal directories ("doug", "joe", etc.) and directories that were the repositories of related modules of source code. Programs that relate to CDS concepts are found in directories "af01", "cr01", and "ed01". Documentation is found in "doc". Files of Job Control Language are in "jcl". Test data is in "test". Directory "r3" contains directories of CDS Release 3 files. The ability to create

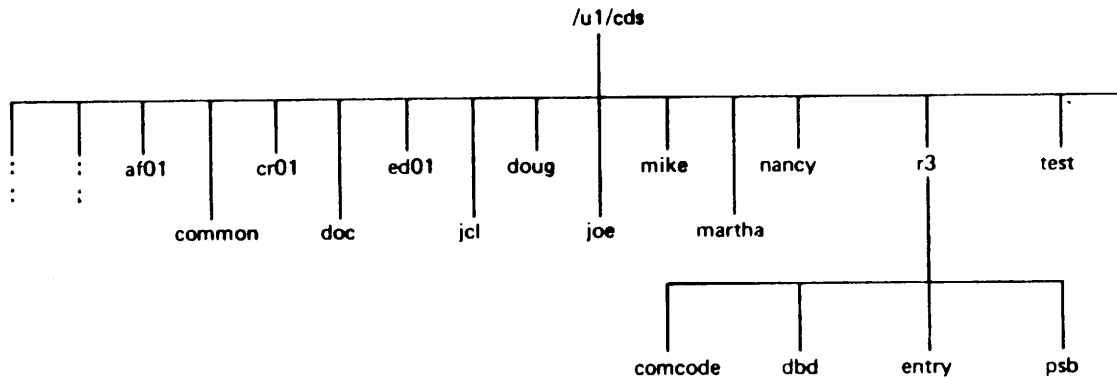


Figure 1. The CDS File System

meaningful collections of data files, reasonably named, under directories, also reasonably named, proved to be a major asset. It was now possible to produce quickly lists of all programs coded, those related to a particular concept, those which adhered to particular naming conventions, and those related to a particular concept that adhered to particular naming conventions.

In theory, this type of grouping and classifying of ideas or programs is possible on other computer systems through naming conventions. But PWB provided the tools that allowed one to go right into the file system and make the computer do the searching. By the time CDS consisted of 140 independently compiled procedures, this feature became invaluable. Our cross-reference listing procedure is shown in detail later.

4.3 Send—the RJE Program

The *send* command allowed us to communicate with the IBM target computer via Remote Job Entry. This is about all that the first version did. The command line

```
send jobcard plijcl source
```

would send the file "jobcard", followed by the file "plijcl", followed by "source". But *send* also assigned special meaning to the "tilde" (~) character. In particular, a line of the form:

```
~ filename
```

read that file as the source of text, and a line of the form:

```
~ _
```

read from the terminal, with a prompt of "input:", for the text. The "~ _", in what came to be called "send-speak", was put to work. Rather than have the programmer enter all of the file names to be sent on the command line, we had *send* prompt for each input item, and lines like "~ plixclg" (for PL/I Optimizer compile, linkage edit and go) and "~ cdsed01" (for the source code for program CDSED01) were the responses.

We used the "~ filename" form embedded in our code to implement the idea of common code, "comcode" for short. Our "comcode" was stored in a directory by that name. Each file contained one program concept. In the majority of cases these were PL/I DECLARE statements of structures that represented data base segments, input formats, and IMS control structures. There were 200 comcode items averaging eleven lines each by the time we were operational. Each module has an average of 4.2 "comcode" references for a saving of 46 lines per procedure.

There was another common directory called "entry" that contained DECLARE statements for each of the external procedures in CDS. By entering the following lines:

```
~/u1/cds/r3/comcode/orderno
~/u1/cds/r3/comcode/bodyin
~/u1/cds/r3/comcode/icferr
~/u1/cds/r3/entry/plitdli
~/u1/cds/r3/entry/prtdate
```

the programmer could reference 48 lines of code that declared

the ORDERNO and BODYIN segments of our ICF (Incomplete Circuit File) data base, the code defining the ON CONDITION ICFERR that handled errors in calls to the ICF, and entry declarations for the external procedures PLITDLI and PRTPDATE.

4.4 The Text Formatter

Documentation support was provided by the text formatter, *roff*. This program made it possible to sit at a terminal and enter a draft document directly, along with an occasional format command for paragraphs, headings, etc. When finished, one asked *roff* to print the document in formatted form. The raw document was entered and edited using the same text editor that was used for entering code, *ed*. Not only did this save having to learn a separate editor for use with the text formatter, but we could easily include sections of code in our documents, and vice versa, without having to retype them.

We found that the job performed by *roff* was acceptable for the documentation and day-to-day business of building a project. We started preparing program documents, such as explanations of CDS error messages, with this program and found it fairly easy to keep them current and available.

Some people question the "waste of time" of typing one's own documents. We feel that for anyone with a moderate amount of typing skill, it takes about the same amount of time to type as to write by hand. Many authors develop their documents at the terminal from a few notes. So there is no time lost, and the result is as good as or better than that from a typing pool.

Roff was even more helpful for large documents with several co-authors. Each author could have an up-to-date and readable copy of the entire document at all times. Our clerk/typist found it more rewarding to be able to correct errors or rearrange paragraphs without having to retype an entire page or "cut and paste," because all the time spent working on the document was productive.

5. RECENT PWB TOOLS

More recently, PWB has increased in potential and CDS has made use of that potential.

The *ed* program has not changed significantly in the last two years, but the few changes have increased the ease with which it can be used to do the more esoteric editing that the experienced user inevitably desires.

Send, however, has grown in capability to the point that it is a major tool in easing the "nuisance work" most programmers have to deal with.

Early modifications to *send* added the ability to establish "keywords" that would prompt the user, who would then respond with appropriate answers. These were substituted into "canned" Job Control Language files, creating custom JCL for the particular purpose at hand. The answers to the keyword prompts would also be displayed prominently in comments, so that if there was

a problem, it was not necessary to try and dig out their values. At the same time we made the prompts for file input request the appropriate type JCL, PL/I source, run data, control cards, etc.

The major benefit achieved here was that it was *impossible* to forget a substitution because you were *always* asked. The result was that we had engineers, technical assistants, and clerks who, by learning the correct responses, repeatedly sent jobs to do testing without ever seeing a JCL statement. Two JCL "gurus" managed everything; JCL syntax errors became so infrequent as to be curiosities. Care was taken to make the prompts for both key words and file input consistent with the *intent* of the job, and not the details of the JCL or IBM file setup. The people thought of *what* they wanted to do; PWB performed the actual work of implementing those intents. By use of these prompts, our non-JCL oriented users could have great flexibility in sending their jobs and still not worry about the details.

Our early JCL files emulated many of the features normally provided by catalogued procedures. However, as our experience increased, we found that we were doing things with our JCL files that were not easily accomplished in catalogued procedures. In particular, one keyword could be used to specify fields anywhere on the JCL statements and in the source text. Also, we could write JCL which referenced other JCL files, thus avoiding duplication and easing maintenance.

5.1 The Source Code Control System

Over the years there have been attempts to provide a means to store, control, and document code as it is being developed. In almost every case these systems incorporate the means of editing the code. The PWB Source Code Control System (SCCS) does not (R0C75A). The programmer requests a copy of the code for editing and SCCS locks out any other edit requests. The programmer then edits the source, which is an ordinary text file, by whatever means available, usually *ed*. When convinced that the new form is what is desired, the user asks SCCS to record the changes and unlock the master file to other editing. SCCS also records a statement of why those changes were made.

SCCS puts no restriction on what the text is or how it is generated and edited. Thus, when CDS started using the SCCS, our programmers had only to learn the initial request, *get*, and the final request to record the changes, *delta*.

5.2 NROFF—the New Text Formatter

Nroff, the new *roff*, has actually been available the entire time CDS has been on the Workbench. However, its much greater power was gained at the cost of syntax and features which are difficult to learn. Recently, the availability of a comprehensive set of "macros" for doing documentation has made *nroff* as easy to use as *roff* (MAS76B). All of the late CDS documentation was done using *nroff*, giving superior document appearance and content in considerably less time than was previously possible.

6. HOW CDS USES PWB

PWB consists of many different processors, some of which perform quite primitive functions. It is the user's responsibility to put these programs together in imaginative and useful ways. It takes a while to get used to the idea that most of the work to accomplish a particular task has been done for you and that your work consists mostly of piecing together the features and little programs you need to produce the desired effect.

For example, when we want to produce a sorted list, we call on *sort* to do it. We do not know or care how it gets its job done or how much machine core or time it takes. We just call it with arguments that indicate what is to be sorted and by what rules. It then goes and does it. The same goes for printing, editing, file searching, string searching, etc. The little programs do some little thing in a reliable and flexible way. We piece them together to do what we want.

Thus the Workbench becomes a human oriented computer system. We spend our time working on what we want to do and how, but on a very high level. The implementation details are not our concern.




Figure 2. The Program Development Sequence

6.1 An Example of PWB Working for CDS

To show just how much work PWB does and how easy it is to get it to do that work, we will trace through a basic terminal session to change a program module, test it via compiles and runs on the target computer, and then make the changes official and permanent. The sequence is shown in Figure 2.

The first step in our example terminal session is to retrieve a copy of the original program module and to have SCCS restrict access to that module to non-editing only. The entered command line is:

```

get -e /u1/cds/cr01/s.ccanal
3.1
106 lines
  
```

That *gets* the SCCS source module "s.ccanal" for editing. The computer's response is printed in **bold**. A file "ccanal" is created for the programmer. SCCS tells us that the current release and level of the module are three and one, respectively, and that the created file has 106 lines.

We will not show the editing process which is fairly standard. Suffice it to say that there are no special considerations that the programmer must make for SCCS while editing the module.

```

send ../jcl/mhb      a job-card file
CLASS=b            "B" class job (sets core and time limits)
JCL:               request for job control cards
- ../jcl/plixc      file for basic compile
RELEASE=3          release of "comcode" to be used
pl/i source:       request for code to be compiled
-ccanal            reference to file to be compiled
pl/i source:       request for further code to be compiled
-                 conclusion of "plixc"
-                 request for further job control cards
JCL:               conclusion of "JCL" prompt and job stream
-                 user information from send
125 cards.
Queued as /u1/hasp/xmit120.
  
```

Figure 3. User Conversation to Compile for Error Messages

When the programmer wishes to compile this program to check for compiler messages, the "conversation" in Figure 3 is held with PWB. What the user does *not* see is nine lines of JCL which include customized comments to help identify this job and 116 line of source code.

The "JCL:" and "pl/i source:" prompts are nested and repeated, allowing multiple compiles in one step and/or multiple steps in a job. The "-" discontinues the current level of nesting.

6.1.1 Compile for Testing. The edit and "plixc" cycle is repeated until the program compiles cleanly. The next step is to put it out where it can be run in a test environment. This is done with the conversation shown in Figure 4. We use the same job card and the prompts have the same meaning. The file "../jcl/compile.test" contains JCL to do a compile and linkage edit into a target machine library, R9411.CDS.R3.TEST.

```

send jcl/mhb      a job card file
CLASS=d          "D" class job (sets core and time limits)
JCL:             request for job control cards
.../jcl/compile_test file to compile for testing
FILE=ccanal     the file to be compiled
PROC=ccanal     the name it will have when released
RELEASE=3       the release it is to be tested in
JCL:             more JCL requested
.../jcl/testlib  file to linkage edit into the test library
DIRECTORY=cr01 the directory of the main procedure
MAIN=cdscr01   the main procedure name
RELEASE=3       to be tested at this release
JCL:             more JCL requested
.../jcl/btsbatch execute the test library
d class job     a reminder that this must be run class "D"
RELEASE=3       the release of the test library to be used
BTS input:      request for input data
.../data/cr01_test input data file
BTS input:      request for input data
...             conclude test data input
JCL:             more JCL requested
...             conclude jobstream
406 cards       user information from send
Queued as /u1/hasp/xmit590.

```

Figure 4. Compile and Linkage Edit for Testing

The linkage editor has the ability to identify the load module with a "stamp" and to provide aliases for the module. In all source files, as a convention, there appear lines of the form:

```

!echo " identify *ccanal1(%R% %L% %D% %T%)" > ccanaLi
!echo " alias ccnext,cccompi" >> ccanaLi

```

The portions within double quotes are linkage edit control cards. These lines, when read by *send*, cause a file "ccanaLi" to be created containing the card images. Then "compile_test" reads that file at the appropriate moment, and the cards become part of the job stream. The "ccanaLi" file is removed later. The "%R%", "%L%", "%D%", and "%T%" are used later for SCCS release, level, date and time. Since we are not using SCCS here, but are only compiling for test, they will not be changed and their presence in the module "stamp" signals that this module is an unofficial version.

The ".../jcl/testlib" file also contains a linkage edit step. It links the main module with all of its supportive modules to create an executable load module in the target machine's dataset called "R9411.CDS.R3.TESTLIB". We also stamp it with the date and time the job originated.

Finally, we add the file in ".../jcl/btsbatch" to the jobstream which exercises the test library with the data in ".../data/cr01_test".

Both "compile_test" and "testlib" reference the shell procedure "auto_compress". This eight line "program" keeps count of how many updates have been sent to the partitioned datasets used by the linkage edit steps. When the count gets to a certain number, currently eight, a job is kicked off that compresses that particular dataset. Since we have begun using this technique we have not had to worry about the problem of doing "garbage collection" on our partitioned datasets. Thus PWB is overcoming a deficiency of the target computer and relieving us of work we really should not have to do. The lines listed below are from "compile_test". They reference the "identify" file (-PROC_i), remove it, and execute the automatic compress. The same type of "send-speak" appears in "testlib".

```

//lked.sysin dd *
- PROCi
-!rm -f PROCi > /dev/null
-!auto_compress test RELEASE; exit 0

```

```

send jcl/cds      job card for administrator
CLASS=b
JCL:
~jcl/compile
DIRECTORY=cr01
PROC=ccanal
RELEASE=3
3.2              new release and level from SCCS
114 lines        new line count from SCCS
JCL:
~jcl/pgmlib      linkage edit into official program library
DIRECTORY=cr01
MAIN=cdscr01
RELEASE=3
JCL:
~
238 cards.
Queued as /u1/hasp/xmit088.

```

Figure 5. Final Compile and Linkage Edit for Release

6.1.2 *Compile for Release.* When the testing is completed, the module must be made "official". We use *delta* to provide protection and keep a history of the changes. The user would type the SCCS command line:

```

delta ../cr01/s.ccanal
history? Change the choice code analysis - tr 76061-1
87 unchanged
27 inserted
19 deleted

```

The final steps are to compile the *deltaed* module and to link it into the "official" executable library. This is done by the CDS program administrator, who sends the job stream shown in Figure 5. The information from SCCS appears because the "jcl/compile" file contains a *get* (without edit) to obtain the PL/I source for send. The command line within that file is:

```

!get /u1/cds/DIRECTORY/s.PROC -rRELEASE -l -p

```

The "-l" asks for the complete history of this file to be put in a file called "l.PROC" (in this case, "l.ccanal") and the "-p" causes the output of *get* to go directly to *send* without using any intervening file.

The %R%, %L%, etc. we saw earlier now come into use. SCCS changes %R% to the release number, %L% to the level, and %D% and %T% to the date and time the *get* was done. This information is scattered through the source code on comments to help the user, and appears on the identify card we saw earlier in the "PROC_i" file, ccanaLi:

```

identify *ccanal1(3.2 76/03/13 17:39:09)
alias ccnext,cccompi

```

So now our load module is stamped with the release, level, date, and time of our module. We now can easily determine if a particular version is up to date.

The "jcl/compile" has an extra step in it to print the history that SCCS provided as part of the compile listing. Since that history tells when and who did what to this module, the listing produced is a complete document of this module to date. That's a handy thing to have, especially if people get into the habit of giving reasonable histories to *delta*. Histories that say "debug" are not all that useful.

The "jcl/pgmlib" file is much like the "jcl/testlib" we saw earlier. However, the read permission is restricted to the CDS program administrator so that only that person is able to send it. It also has an extra step after the linkage edit to create a listing of all the "identify" stamps that we put on our load modules. Thus the linkage edit listing includes a complete list of all CDS modules in the executable module, including their SCCS release

level, and the date and time they were retrieved from SCCS, which was a useful thing to have when we were not sure what version was last compiled.

Again, "jcl/compile" and "jcl/pgmlib" both use the "auto_compress" to keep things on the target machine clean.

We would point out that most of the automatic processes we show would have to be performed by hand or by writing special programs if we were to use the facilities available on the target computer. PWB has totally relieved us of the drudgery of the manual process and even hidden the work being done.

6.2 Using PWB To Analyze Output

The RJE process permits the returning of output to PWB instead of having it printed. The "big file scanner", *bfs*, is used to scan large files. The authors have used this to some advantage. We have a *bfs* script that searches PL/I compiler output for the significant diagnostics, linkage editor complaints, and the printout from the actual run. When the project goes into "panic mode," this is very useful for first compiles and test case drivers. We also use it on a casual basis during the normal work day. Being able to look at twenty lines of significant diagnostics rather than twenty pages of output is convenient and we still can go back for details.

7. IS PWB WORTH IT?

That is the real point of all this: is it really worth learning another system to gain the benefits of PWB?

7.1 Productivity

PWB increases programmer productivity in a number of significant ways.

7.1.1 Fewer Steps in Coding. Generally, one step is eliminated in transferring an idea into the code of a program. The usual sequence of events for a batch card-oriented system is:

- Rough draft the idea into a flowchart, or some code or shorthand form.
- Expand into code on a coding form, hand written.
- Key punch the code (either by the author or by a keypunch service).
- Wrap the code in a JCL deck and take to the computer center to be compiled.

In CDS we have observed that the second step is frequently skipped. The programmer arrives at the terminal with a rough draft of what is intended and refines it while entering it via the editor. So the terminal serves the purpose of the coding form.

Obviously, there is no wait for the keypunch service nor are any physical cards generated. And finally there is no need to go to the computer counter to push the deck across. The *send* command does that for you.

When that first compile comes back with its inevitable diagnostic messages, the next savings are realized. The programmer can log into UNIX and directly add that missing comma, include the forgotten argument, or move a misplaced statement. There is no need to write up another coding sheet, or duplicate cards, or shuffle cards. Just log in, correct, and *send* again.

Thus, in terms of the productivity to be gained through reductions in duplicated effort and trips to the computer center, PWB provides significant enhancements relative to the card-oriented environment.

7.1.2 "Automatic" Documentation and File Maintenance. Since the Source Code Control System keeps both the code and the history of updates, and since the Job Control files are flexible and yet always consistent, the Workbench performs all the functions normally assigned to a Program Librarian. We have a "daily daemon" which runs every weekday at 5 am. It "mails" reports to programmers on files added and removed since the previous

working day. If something disappears, either through a system failure (which is rare), or a programmer error (much more common), we usually know about it within 48 hours and can get it backed up. On Fridays, it sends a job to produce a usage report on all of our IBM datasets so that we can stay ahead of our requirements, rather than reacting to crises. Before we built the "auto_compress" discussed in Section 6.1.1, the daily daemon also sent a job to compress all of our partitioned datasets.

In terms of the productivity gained by automating the work associated with staying ahead of the demands that CDS was making on the target machine, PWB was again very helpful.

7.1.3 Non-Programmer Productivity. As we mentioned earlier, there was a fair amount of telephone engineering being done in the CDS project. The people who were involved did not know, and did not wish to learn, the various intricacies of JCL and IMS which the programmers live with. By using the Workbench as a filter, they were not forced to learn these extraneous systems, and thus could concentrate on designing CDS.

The programmers and those of us providing the JCL and IMS support also found the layer of filtering helped us concentrate on getting the programs working. Far less time was spent chasing down JCL syntax errors, recovering from dropped or misplaced decks, or counseling people on how to read cryptic messages.

True, there was a price to pay: learning enough about the key features of the PWB to make it work for us. But relative to the extraneous education we avoided for eight of our people, the time spent learning about UNIX was not very significant.

7.2 Better Code?

It is one thing to do a job faster. Does PWB help to do it better?

7.2.1 Program Style and Structure. The popular concepts of program structure and style are much touted in the literature, but we suspect most software shops are finding them difficult to implement. It is just a pain to have to re-code working code, "just to make it look pretty". We found that the use of PWB helped and even encouraged our programmers to write new code using the "good style" concepts. It was also possible to take existing code and "structure" it without changing a single character of actual code. One simply spaced it out, indented, and blocked as required. That did not change anything as far as the final compiled machine code was concerned, but made it *much* easier to maintain.

The ability to plagiarize well written code and modify it just a little bit was quickly discovered. This technique was used when a piece of common code could not be constructed for a particular purpose. A generalized solution to the problem would be made available and each programmer would adapt it as needed.

7.2.2 Sharing Code. The "comcode" idea mentioned earlier provided consistency in naming and usage of CDS concepts everywhere they appeared, and they appeared everywhere. Programmers working on opposite ends of the system had no trouble talking about data concepts that they had to share since they shared the same "comcode" for those concepts. The first time that inter-program communication was attempted via our data base, it worked! Also, since our "comcode" items were liberally and intelligently commented, every program that used them benefited.

Just as important was the effort saved. On the average, the programmer saved ten lines of coding every time a comcode was referenced. Thus, the naming conventions were easy to enforce since it was easier to use them than not.

7.2.3 Keeping Names Meaningful. Many more times than once in CDS we were faced with this problem: This variable no longer means what its name implies. It should really be changed to be more meaningful. "But it is used all over the place! How can you be sure you've gotten every occurrence?" Normally that is

a sticky problem. It is even worse when it is not just the name, but the nature of the data it represents that changes.

With the PWB we could, and did, find every occurrence of a variable and change its name and nature, without causing the usual catastrophies. To change the size and structure of the root segment of our principal data base in 100 independent PL/I procedures took one evening's worth of work for the authors. We were up and running the next morning. As we recall, there were three bugs associated with that change, all due to an oversight on our part, and they were all found and cleared by the end of the week. The code to accomplish this is shown in Section 7.3.3.

7.3 Because It's Only Impossible, We'll Give You Until Tomorrow

In any job there are those aspects which can best be described as a nuisance. And yet they take up time and effort and must be dealt with. Sometimes, they are extremely difficult to do, other times just plain tedious. It was pleasantly surprising to find that the Workbench could relieve the tedium and ease the difficulties.

7.3.1 How Many Statements in CDS? Programmers will never understand the fascination that the managers of software projects have with "statements of code." Normally when faced with the request for a count, programmers cringe, as did we. But when pressured we discovered that it was not all that bad. We found there were already programs for extracting lines containing particular characters (*grep*) and counting lines (*wc*) and we could pipe the output of one into the other. We built a shell procedure that looked at all of our code and counted semi-colons, the statement delimiter in PL/I. Thus the Workbench reduced a tedious job to a trivial one.

7.3.2 Where Did That Extra Character Come From? In any job, no matter how hard you try to avoid it, something that you did not anticipate takes almost as much time as planned activities. In CDS, the terminal we selected provided the "diversion." Simply stated, the relationships between the cassette tape drive and the data line to the computer were not as advertised. Characters which were recorded on the tape would not get to the computer. Characters which were not on the tape would be transmitted. Since these were non-graphic and control characters, we had some difficulty isolating the problem.

But UNIX came to our rescue. By turning off the special meaning that characters had to UNIX, that is by setting the input processor to "raw" mode which is within the user's power, we were able to record exactly what the terminal sent down the line in a file. The *od* program, octal dump, was then used to see exactly what was present. Through this technique we were able to determine the specific terminal deficiencies that were causing our problems. We then provided the manufacturer the specification for the correction which we required.¹

This discussion brings up an important point. We have used several brands of printing and CRT terminals. In each case we could tailor the character set and response times from UNIX so that the terminal could keep up. If the terminal is fast, as in the CRT types, UNIX can be made faster to save time. It should be understood that we are not talking about changes to system code; we mean that the user enters a simple "set teletype" command (*stty*) to provide the features required.

7.3.3 Change the World. We talked earlier about finding and making massive changes reliably. That process is not trivial. It requires a thorough understanding of what needs to be done. However, given that understanding, PWB provides the tools to reliably search out and change what is needed. The keystone of those tools is the ability to create files of commands to drive the

text editor. Moreover, it is a simple matter to make the files dynamic, changing according to need. All of this is done at the command language level, and is readily learned if one is willing to spend the time to read the programmer's manual and experiment a little.

For example, the need came up recently to change all occurrences of the character string "%M%" to the name of the file in which that string was found. There were 130 files spread around nine directories. The problem was to create the commands to *get* the file out of SCCS, edit it, and *delta* the result back in. The following was the heart of the solution.

```
if ! { get -e $1 } exit
gath -s III=$1 - ; delta $1 "-ychange %M% to proc name"
- $ed III.a > /dev/tty ; exit 0
- $f
- $g/%M%/s//III/gp
- $w
- $q
```

What it does is:

- When called with the file name as argument (\$1), does a *get* of the file for editing (-e) and if that fails, exits.
- *Gath* establishes the string "III" as the file name and reads the lines that follow (-).
- The "--\$" lines are read by *gath* and "III" is replaced with the file name. The resultant lines are given to the shell (*sh*) and thus the *ed* command performs the "f", "g", "w", and "q" commands which edit the file.
- The *delta* is made on the file with the given history.

A driver already existed for finding all files in the needed directories. It took about a half hour to develop and test the above code and about 50 minutes during off-hours for the PWB to perform the actual work. We've kept the skeleton around, since this type of thing comes up about once every other month.

7.3.4 Compile the World. In *send*, the keyword capability is not limited just to prompting. The few times it became necessary to compile everything we had we would satisfy the keywords in advance, using files of definitions or shell procedures, and then use the same JCL files as we did every day to do the work. The result was 100% consistency with day-to-day practice.

```
rm -f /u1/cds/xrefi
: loop
if $!x = x goto wrapup
chdir /u1/cds/$!
get . -r3 -s
grep "- " ?[0-z]* >> /u1/cds/xrefi
rm -f ?[0-z]*
shift
goto loop
: wrapup
chdir /u1/cds/r3/comcode
grep "- " * >> /u1/cds/xrefi
chdir /u1/cds/r3/entry
grep "- " * >> /u1/cds/xrefi
ed - /u1/cds/xrefi
l,$s://
w
q
reform +16 +8 < /u1/cds/xrefi | sort +1 > /u1/cds/xrefs
rm /u1/cds/xrefi
mail jw
xref completed
```

Figure 6. Shell Procedure for Cross-Referencing Comcode

¹ The final result of the investigation was that line-feed characters recorded on the tape cassette did not get sent to the computer and every block from the tape that was transmitted terminated with two carriage-return characters. IMS teleprocessing required the line-feed character and could not tolerate the extra unwanted carriage-returns.

7.3.5 Cross-Reference the World. The Programmer's Workbench gave us the tools that allowed us to build a shell procedure to find all occurrences of the "tilde references" and tell us the

program module they occurred in sorted by the order of the comcode files. Time from conception to tested procedure was one day. We list it in Figure 6 to show its size. Briefly, it:

- Removes the temporary file xrefi.
- From "": loop" to "goto loop" gets all programs in the directories specified. All the lines which have "-" in them are put into "xrefi" along with the file name.
- From "": wrapup" to "grep "-" * >> /u1/cds/xrefi" does the same for the "comcode" and "entry" directories.
- From "ed" to "reform" produces a formatted, sorted list of all the references and puts it into "xrefs".
- Remove the temporary file "xrefi".
- Mail to the login "jw" the message "xref completed".

It is a bit complicated. But this was created when we were still novices. A newer version, which produces identical output, is faster, more robust, and the user does not have to provide the list of directories to be used.

The "xref" has proven valuable time and time again, when a CDS system concept was changing and we wished to see what modules were impacted. More importantly, we do not feel it would have been feasible, with the development schedule we had, to build a similar system for use with LIBRARIAN or any of the other tools available on the IBM machine. Certainly, it would have been difficult to do it in one day.

8. DEFICIENCIES

8.1 Education

When the authors first started using the PWB we were told that our login code would be "cds" and were handed a copy of the UNIX Programmer's Manual. This eight part volume was all the education that was initially provided. It soon came to be known as "The Book", a phrase taken from the expression often heard in our office, "It's in The Book." Truly all the information required to properly use the system was in the book but in an incredibly terse format intended for reference only. The entire writeup on *ed* consists of four pages. A part of the problem is that in almost any software project, PWB included, the paperwork comes after getting the product out the door.

Recently the user community has expanded to such an extent that the PWB developers do not have enough time to answer all the user's questions individually. This has led to a series of useful tutorials, memoranda, and classes.

8.2 File System Management

As we have seen in our example, use of the full path name of a file aids in adding robustness to the code. The shell procedure or tilde reference will always work. However, the amount of space that exists in a file system such as "/u1" is limited and sometimes it becomes necessary to move a project to another file system. Currently, there is no support for such a move, and we were stuck with the job of finding all occurrences of "/u1" and changing them to "/u9". It was not all that difficult to do, but the PWB "super-user" (system administrator) could have done it for us in much less time.

In PWB all storage is on-line. Although daily backups are taken, there is no convenient way for users to archive old material. Source code files in SCCS format can have a lot of built-up fat. Old releases which are not currently in the field could be "crunched" out from the bottom if there were some way to store the historic copy offline.

8.3 Independent PWB versus On-Line

For most of the users of PWB, it is an interface to a large target machine. This immediately brings up the objection that PWB is no better than being in a batch environment. There is no general capability for on-line compilation and testing. These func-

tions might be gained if the target provided a time sharing system, such as IBM's Time Sharing Option (TSO).

But would it be worth while to build a PWB facility under TSO (or other target operating system)? Certainly the ability to do on-line compilation and testing would exist. This approach may be appropriate at some installations. The general issues involved in using an independent PWB are discussed at length in [DOL76A].

If CDS were developed on a TSO Workbench, we would be sharing the target machine with a large batch environment and an on-line IMS environment which is already having difficulty servicing its terminals adequately. With PWB serving a strictly on-line environment, and the vast majority of its customers using its editing, RJE, and documentation facilities, we feel it is doing a better than adequate job. From our particular experience with CDS, the lack of on-line compilation and testing seems more than compensated for by the other advantages of the PWB. Also, in our installation (as in others possessing finite resources), we have often seen contention for priority of use between batch and on-line applications, and we welcome PWB's ability to provide service without getting involved in such contention.

9. RELIABILITY

The reliability of the PWB system has been very good. In over two years of use, only one file of CDS source code was damaged. It was restored to the previous day's version an hour later.

File space availability is a recurring problem. Unlike other systems in which users have a fixed allocation, PWB shares the available space across a group of users. Some users are not as nice about staying under their "paper" allocations as others: about every six months a file system runs out of space.

Because our location currently has four PWB machines, if any one malfunctions files can be switched to an alternate machine giving a degree of "fail-soft." Thus even when a machine is broken the users have service, though admittedly degraded. Even so we have found that the system uptime is very good. It is generally available during working hours, although at about 2:00 pm access can be limited due to line congestion. Usually it will stay up throughout an entire weekend without any attendant.

10. CONCLUSIONS

We have shown a small sample of the experiences of one group using the PWB to aid in developing a real-life application.

Since we first joined, the Programmer's Workbench community has grown at an astonishing rate. It is being used by development projects and maintenance projects, documentation centers and typing pools, clerks, typists, programmers, engineers, and supervisors. It is probably safe to assume that no two use it in exactly the same way.


But there is little doubt that the availability and capabilities of the UNIX Programmer's Workbench are having major impact wherever it is used. It reliably provides useful computing power to a large and diverse community at a very low cost. Moreover, this power is available in the form most people want it: a human oriented system which is easier to use than not.


This type of software design, where the system does not drive the user, but rather the user easily drives the system, will, in our opinion, have great and favorable impact on professional and public acceptance of future computer technology.

REFERENCES

All references cited in this paper appear at the end of "An Introduction to the Programmer's Workbench," by Dolotta, T. A., and Mashey, J. R., in these Proceedings.







PWB/UNIX Documentation Roadmap

J. R. Mashey

Bell Laboratories
Piscataway, New Jersey 08854

1. INTRODUCTION

A great deal of documentation exists for PWB/UNIX. It has different formats, is contributed by many different people, and is modified frequently. New users are often overcome by the volume and distributed nature of the documentation. This "roadmap" attempts to be a terse, up-to-date outline of crucial documents and information sources.

Numerous people have contributed comments and information for this "roadmap," in order to make it as helpful as possible for PWB/UNIX users. *However, many of these comments are accurate only with regard to PWB/UNIX and may well be totally inapplicable to other versions of UNIX.*

1.1 Things to Do

See a local PWB/UNIX system administrator to obtain a "login name" and get other appropriate system information.


1.2 Notation Used in this Roadmap

- {N} → Section N in this "roadmap."
- ++ → item required for everyone.
- + → item recommended for most users.

All other items are optional and depend on specific interest (a list of relevant documents appears in the Table of Contents of *Documents for the PWB/UNIX Time-Sharing System*).

Items in Section N of the *PWB/UNIX User's Manual* are referred to by *name(N)*.

1.3 Prerequisite Structure of Following Sections



2. BASIC INFORMATION

Don't do anything else until you have learned most of this section. You must know how to log onto the system, make your terminal work correctly, enter and edit files, and perform basic operations on directories and files.

2.1 PWB/UNIX User's Manual ++

- Read *Introduction* and *How to Get Started*.
- Look through Section I to become familiar with command names.
- Note the existence of the *Table of Contents* and of the *Permuted Index*.

Section I will be especially needed for reference use.

2.2 UNIX for Beginners ++

2.3 A Tutorial Introduction to the UNIX Text Editor ++

2.4 Advanced Editing on UNIX +

2.5 PWB Papers from the Second International Conference on Software Engineering +

Gives an overview of the Programmer's Workbench.

2.6 Things to Do

- Do all the exercises found in {2.2} and {2.3}, and maybe {2.4}.
- Create a file named “.mail” in your login directory,* so that other people (such as system administrators) can send you mail. This can be done by:

```
cp /dev/null .mail
```

- If you want some sequence of commands to be executed each time you log in, create a file named “.profile” in your login directory containing the commands you want executed. For more information, see *Initialization* in *sh(1)*.
- Files in directory “/usr/news” contain recent information on various topics. To see what has been updated recently, type:

```
ls -lt /usr/news
```

and then print any files that look interesting. Other useful actions include:

mail -f /usr/news/.mail	gives recent history from primary system mailbox.
cat /usr/news/helpers	gives contacts and telephone numbers for counseling, file restorals, trouble reporting, and other services.
nroff -mm /usr/news/roadmap	prints current copy of this “roadmap.”
cat /usr/news/terminals	gives recommendations on selection of computer terminals.

2.7 Manual Pages to Be Studied

The following commands are described in Section I of the *PWB/UNIX User's Manual*, and are used for creating, editing, moving (i.e., renaming), and removing files:

cat(1)	concatenate and print files (no pagination).
chdir(1)	change working (current) directory; a.k.a. <i>cd(1)</i> .
cp(1)	make a copy of an existing file.
ed(1)	text editor.
ls(1)	list a directory; file names beginning with “.” are not listed unless the “-a” flag is used.
mkdir(1)	make a (new) directory.
mv(1)	move (rename) file.
pr(1)	print files (paginated listings).
rm(1)	remove (delete) file(s).
rmdir(1)	remove directory(ies).

The following help you communicate with other users, make proper use of different kinds of terminals, and print manual pages on-line:

login(1)	sign on.
mail(1)	send mail to other users; inspect mail from them, or contents of the system mailbox.
man(1)	print pages of <i>PWB/UNIX User's Manual</i> .
stty(1)	set terminal options; i.e., inform the system about the hardware characteristics of your terminal.
tabs(1)	set tab stops on your terminal.
terminals(VII)	gives descriptions of commonly-used terminals.
who(1)	print list of users currently logged in.
write(1)	communicate with another (logged in) user.

* The directory you are in right after logging into the system.

3. BASIC TEXT PROCESSING AND DOCUMENT PREPARATION

You should read this section if you want to *use* existing text processing tools to write letters, memoranda, manuals, etc.

3.1 PWB/MM—Programmer's Workbench Memorandum Macros ++

This is a reference manual, and can be moderately heavy going for a beginner. Try out some of the examples, and stick close to the default options.

3.2 Typing Documents with PWB/MM ++

A handy fold-out.

3.3 NROFF/TROFF User's Manual +

Describes the text formatting language in great detail; look at the REQUEST SUMMARY, but don't try to digest the whole manual on first reading.

3.4 Documentation Tools and Techniques +

This overview of UNIX text processing methods is one of the papers from the Second International Conference on Software Engineering. (See {2.5} above).

3.5 Manual Pages to Be Studied

mm(I)	makes it easy to specify standard options to <i>nroff</i> (I).
nroff(I)	read to see formatter option flags.
spell(I)	identifies possible spelling errors.
tmac.name(VII)	list of text-formatting macro packages.
typo(I)	identifies possible typographical errors.

To obtain some special functions (e.g., reverse paper motion, subscripts, superscripts), you must either indicate the terminal type to *nroff* or post-process *nroff* output through one of the following:

450(I)	newer Diablo printer terminals, such as the DASI450, DIABLO 1620, XEROX 1700, etc.
col(I)	terminals lacking physical reverse motion, such as the Texas Instrument 700 series.
gsi(I)	older Diablo printer terminals, such as the GSI300, DASI300, DTC300, etc.
hp(I)	Hewlett-Packard 2640 terminals (HP2640A, HP2640B, HP2644A, HP2645A, etc.).

4. SPECIALIZED TEXT PROCESSING

The tools listed in this section are of a more specialized nature than those in {3}.

4.1 TBL—A Program to Format Tables +

Great help in formatting tabular data (see *tbl*(I)).

4.2 Typesetting Mathematics—User's Guide (2nd Edition) +

Read this if you need to produce mathematical equations. It describes the use of the equation setting commands *eqn*(I) and *neqn*(I).

4.3 A TROFF Tutorial

An introduction to formatting text with the phototypesetter.

4.4 Manual Pages to Be Studied

diffmark(I)	marks changes between versions of a file, using output of <i>diff</i> (I) to produce "revision bars" in the right margin.
eqn(I)	preprocessor for mathematical equations (phototypesetter).
neqn(I)	preprocessor for mathematical equations (terminals).
tbl(I)	preprocessor for tabular data.
troff(I)	formatter for phototypesetter.

5. ADVANCED TEXT PROCESSING

You should read this section if you need to *design* your own package of formatting macros or perform other actions beyond the capabilities of existing tools; {3} is a prerequisite, and familiarity with {4} is very helpful, as is an experienced advisor. It takes a great deal of effort to write a good package of macros for general use. Don't reinvent what you can borrow from an existing package (such as PWB/MM).

5.1 NROFF/TROFF User's Manual ++

Look at this in detail and try modifying the examples. If you are going to use the phototypesetter, do the same for *A TROFF Tutorial* ({4.3} above).

5.2 Things to Do

It is fairly easy to use the text formatters for simple purposes. A typical application is that of writing simple macros that print standard headings in order to eliminate repetitive keying of such headings. It is extremely difficult to set up general-purpose macro packages for use by large numbers of people. If possible, try to use an existing package or modify one as needed. Look at existing packages first—see *mac.name*(VII).

5.3 Manual Pages to Be Studied

All pages mentioned in {3} and {4}.

6. COMMAND LANGUAGE (SHELL) PROGRAMMING

The Shell provides a powerful programming language for combining existing commands. This section should be especially useful to those who want to automate manual procedures and build data bases.

6.1 The UNIX Time-Sharing System ++

6.2 PWB/UNIX Shell Tutorial ++

6.3 Things to Do

If you want to create your own library of commands, create a “.path” file in your login directory, as described in *sh*(1).

6.4 Manual Pages to Be Studied

Read *sh*(1) first; the following pages give further details on commands that are most frequently used within command language programs:

<i>echo</i> (1)	echo arguments (typically to terminal).
<i>equals</i> (1)	Shell assignment command (for variables).
<i>exit</i> (1)	terminate command file.
<i>expr</i> (1)	evaluate an algebraic expression.
<i>fd2</i> (1)	redirect diagnostic output.
<i>if</i> (1)	conditional command.
<i>next</i> (1)	read command input from named file.
<i>nohup</i> (1)	run a command immune to communications line hang-up.
<i>onintr</i> (1)	handle interrupts in Shell files.
<i>pump</i> (1)	Shell data transfer command.
<i>sh</i> (1)	Shell (command interpreter).
<i>shift</i> (1)	adjust Shell arguments.
<i>switch</i> (1)	Shell multi-way branch command.
<i>while</i> (1)	Shell iteration command.

7. FILE MANIPULATION

In addition to the basic commands of {2}, many UNIX commands exist to perform various kinds of file manipulation. Small data bases can often be managed quite simply, by combining text processing (from {5}), command language programming {6}, and commands listed below in {7.2}.

7.1 Things to Do

This "roadmap" notes only the most frequently used commands. It is wise to scan Section I of the *PWBUNIX User's Manual* periodically—you will often discover new uses for commands.

7.2 Manual Pages to Be Studied

The following are used to search or edit files in a single pass:

grep(I)	search a file for a pattern; more powerful and specialized versions include <i>egrep(I)</i> , <i>fgrep(I)</i> , and <i>rgrep(I)</i> .
sed(I)	stream editor.
tr(I)	transliterate (substitute or delete specified characters).

The following compare files in different ways:

cmp(I)	compare files (byte by byte).
comm(I)	print lines common to two files, or lines that appear in only one of the two files.
diff(I)	differential file comparator (minimal editing for conversion).

The following combine files and/or split them apart:

ar(I)	archiver and library maintainer.
cpio(I)	general file copying and archiving.
csplit(I)	split file by context.
split(I)	split file into chunks of specified size.

These commands interrogate files and print information about them:

file(I)	determine file type (best guess).
od(I)	octal dump (and other kinds also).
wc(I)	word (and line) count.

Miscellaneous commands:

find(I)	search directory structure for specified kinds of files.
gath(I)	gather real and virtual files; alias for <i>send(I)</i> .
help(I)	ask for help about a specific error message.
reform(I)	reformat "tabbed" files (often used to truncate lines).
sort(I)	sort or merge files.
tee(I)	copy single input to several output files.
uniq(I)	report repeated lines in a file, or obtain unique ones.

8. C PROGRAMMING

Try to use existing tools first, before writing C programs at all.

8.1 Programming in C—A Tutorial ++

Read; try the examples.

8.2 C Reference Manual ++

Terse but complete reference manual.

8.3 A New Input-Output Package +

Describes a new I/O package that is superseding many of the existing routines; write any new code using this package.

8.4 UNIX Programming +**8.5 YACC—Yet Another Compiler Compiler****8.6 LEX—Lexical Analyzer Generator****8.7 Make—A Program for Maintaining Computer Programs****8.8 Things to Do**

The best way to learn C is to look at the source code of existing programs, especially ones whose functions are well known to you. Much code can be found in directory “/sys/source”. In particular, directories “s1” and “s2” contain the source for most of the commands. Also, investigate directory “/usr/include”.

8.9 Manual Pages to Be Studied

adb(I)	C debugger; more powerful (but more complex) than the older <i>cdb(I)</i> .
cc(I)	C compiler.
cdb(I)	C debugger (for post-mortem core dumps and other debugging).
ld(I)	loader (you must know about some of its flags).
lex(I)	generate lexical analyzers.
make(I)	automate program regeneration procedures.
nm(I)	print name (i.e., symbol) list.
prof(I)	display profile data (used for program optimization).
regcmp(I)	compile regular expression.
strip(I)	remove symbols and relocation bits from executable file.
time(I)	time a command.
yacc(I)	parser generator.

9. IBM REMOTE JOB ENTRY (RJE)

This section is for those who use PWB/UNIX to submit jobs to remote computers.

9.1 Guide to IBM Remote Job Entry for PWB/UNIX Users +**9.2 Manual Pages to Be Studied**

bfs(I)	big file scanner (scans RJE output).
csplit(I)	split file by context (often used to split RJE output).
fspec(V)	format specification in text files.
reform(I)	reformat files (often used to convert source programs from non-UNIX systems).
rjstat(I)	RJE status and enquiries.
send(I)	submit RJE job.

10. SOURCE CODE CONTROL SYSTEM (SCCS)

SCCS can be used to maintain, control, and identify files of text as they are modified and updated. Its most common use is for maintaining source programs, as well as for keeping track of successive versions of various documents; in combination with *diffmark(I)*, this allows one to automatically generate “revision bars” in successive editions of such documents.

10.1 SCCS/PWB User's Manual ++**10.2 Manual Pages to Be Studied**

Of the following, *get(I)*, *delta(I)*, and *prt(I)* are most frequently used.

admin(I)	administer SCCS files (including creation thereof).
chghist(I)	change the history entry of an SCCS delta.
comb(I)	combine SCCS deltas.
delta(I)	make an SCCS delta (a permanent record of editing changes).
get(I)	get a version of an SCCS file.
prt(I)	print SCCS file.

rm δ (I)	remove a delta from an SCCS file.
sccsdiff(I)	get the differences between two SCCS deltas.
what(I)	find and print SCCS identifications in files.

11. NUMERICAL COMPUTATION

11.1 DC—An Interactive Desk Calculator


11.2 BC—An Arbitrary Precision Desk Calculator Language


11.3 RATFOR—A Preprocessor for a Rational Fortran

11.4 Manual Pages to Be Studied

bas(I)	BASIC interpreter.
bc(I)	interactive language, acts as front end for <i>dc</i> (I)
dc(I)	desk calculator.
fc(I)	Fortran compiler/interpreter.
rc(I)	RATFOR preprocessor.









The PWB/UNIX* document entitled:
PWB/UNIX Beginner's Course
is not yet available.

* UNIX is a Trademark/Service Mark of the Bell System.







A Tutorial Introduction to the UNIX Text Editor

B. W. Kernighan

Bell Laboratories, Murray Hill, N. J.

Introduction

Ed is a "text editor", that is, an interactive program for creating and modifying "text", using directions provided by a user at a terminal. The text is often a document like this one, or a program or perhaps data for a program.

This introduction is meant to simplify learning *ed*. The recommended way to learn *ed* is to read this document, simultaneously using *ed* to follow the examples, then to read the description in section I of the UNIX manual, all the while experimenting with *ed*. (Solicitation of advice from experienced users is also useful.)

Do the exercises! They cover material not completely discussed in the actual text. An appendix summarizes the commands.

Disclaimer

This is an introduction and a tutorial. For this reason, no attempt is made to cover more than a part of the facilities that *ed* offers (although this fraction includes the most useful and frequently used parts). Also, there is not enough space to explain basic UNIX procedures. We will assume that you know how to log on to UNIX, and that you have at least a vague understanding of what a file is.

You must also know what character to type as the end-of-line on your particular terminal. This is a "newline" on Model 37 Teletypes, and "return" on most others. Throughout, we will refer to this character, whatever it is, as "newline"

Getting Started

We'll assume that you have logged in to UNIX and it has just said "%". The easiest way to get *ed* is to type

```
ed      (followed by a newline)
```

You are now ready to go — *ed* is waiting for you to tell it what to do.

Creating Text — the Append command "a"

As our first problem, suppose we want to create some text starting from scratch. Perhaps we are typing the very first draft of a paper; clearly it will have to start somewhere, and undergo modifications later. This section will show how to get some text in, just to get started. Later we'll talk about how to change it.

When *ed* is first started, it is rather like working with a blank piece of paper — there is no text or information present. This must be supplied by the person using *ed*; it is usually done by typing in the text, or by reading it into *ed* from a file. We will start by typing in some text, and return shortly to how to read files.

First a bit of terminology. In *ed* jargon, the text being worked on is said to be "kept in a buffer." Think of the buffer as a work space, if you like, or simply as the information that you are going to be editing. In effect the buffer is like the piece of paper, on which we will write things, then change some of them, and finally file the whole thing away for another day.

The user tells *ed* what to do to his text by typing instructions called "commands." Most commands consist of a single letter, which must be typed in lower case. Each command is typed on a separate line. (Sometimes the command is preceded by information about what line or lines of text are to be affected — we will discuss these shortly.) *Ed* makes no response to most commands — there is no prompting or typing of messages like "ready". (This silence is preferred by experienced users, but sometimes a hangup for beginners.)

The first command is *append*, written as the letter

```
a
```

all by itself. It means "append (or add) text lines to the buffer, as I type them in." Appending is rather like writing fresh material on a piece of paper.

So to enter lines of text into the buffer, we just type an "a" followed by a newline, followed by the lines of text we want, like this:

```
a
Now is the time
for all good men
to come to the aid of their party.
```

The only way to stop appending is to type a line that contains only a period. The "." is used to tell *ed* that we have finished appending. (Even experienced users forget that terminating "." sometimes. If *ed* seems to be ignoring you, type an extra line with just "." on it. You may then find you've added some garbage lines to your text, which you'll have to take out later.)

After the append command has been done, the buffer will contain the three lines

```
Now is the time
for all good men
to come to the aid of their party.
```

The "a" and "." aren't there, because they are not text.

To add more text to what we already have, just issue another "a" command, and continue typing.

Error Messages - "?"

If at any time you make an error in the commands you type to *ed*, it will tell you by typing

```
?
```

This is about as cryptic as it can be, but with practice, you can usually figure out how you goofed.

Writing text out as a file - the Write command "w"

It's likely that we'll want to save our text for later use. To write out the contents of the buffer onto a file, we use the *write* command

```
w
```

followed by the filename we want to write on. This will copy the buffer's contents onto the specified file (destroying any previous information on the file). To save the text on a file named "junk", for example, type

```
w junk
```

Leave a space between "w" and the file name. *Ed* will respond by printing the number of characters it wrote out. In our case, *ed* would respond with

68

(Remember that blanks and the newline character at the end of each line are included in the character count.) Writing a file just makes a copy of the text - the buffer's contents are not disturbed, so we can go on adding lines to it. This is an important point. *Ed* at all times works on a copy of a file, not the file itself. No change in the contents of a file takes place until you give a "w" command. (Writing out the text onto a file from time to time as it is being created is a good idea, since if the system crashes or if you make some horrible mistake, you will lose all the text in the buffer but any text that was written onto a file is relatively safe.)

Leaving ed - the Quit command "q"

To terminate a session with *ed*, save the text you're working on by writing it onto a file using the "w" command, and then type the command

```
q
```

which stands for *quit*. The system will respond with "%". At this point your buffer vanishes, with all its text, which is why you want to write it out before quitting.

Exercise 1:

Enter *ed* and create some text using

```
a
... text ...
```

Write it out using "w". Then leave *ed* with the "q" command, and print the file, to see that everything worked. (To print a file, say

```
pr filename
```

or

```
cat filename
```

in response to "%". Try both.)

Reading text from a file - the Edit command "e"

A common way to get text into the buffer is to read it from a file in the file system. This is what you do to edit text that you saved with the "w" command in a previous session. The *edit* command "e" fetches the entire contents of a file into the buffer. So if we had saved the three lines "Now is the time", etc., with a "w" command in an earlier session, the *ed* command

```
e junk
```

would fetch the entire contents of the file "junk" into the buffer, and respond

which is the number of characters in "junk". If anything was already in the buffer, it is deleted first.

If we use the "e" command to read a file into the buffer, then we need not use a file name after a subsequent "w" command; *ed* remembers the last file name used in an "e" command, and "w" will write on this file. Thus a common way to operate is

```
ed
e file
[editing session]
w
q
```

You can find out at any time what file name *ed* is remembering by typing the *file* command "f". In our case, if we typed

```
f
ed would reply
junk
```

Reading text from a file — the Read command "r"

Sometimes we want to read a file into the buffer without destroying anything that is already there. This is done by the *read* command "r". The command

```
r junk
```

will read the file "junk" into the buffer; it adds it to the end of whatever is already in the buffer. So if we do a read after an edit:

```
e junk
r junk
```

the buffer will contain *two* copies of the text (six lines).

```
Now is the time
for all good men
to come to the aid of their party.
Now is the time
for all good men
to come to the aid of their party.
```

Like the "w" and "e" commands, "r" prints the number of characters read in, after the reading operation is complete.

Generally speaking, "r" is much less used than "e".

Exercise 2:

Experiment with the "e" command — try reading and printing various files. You may get an error "?", typically because you spelled the file name wrong. Try alternately reading and appending to see that they work similarly. Verify that

```
ed filename
```

is exactly equivalent to

```
ed
e filename
```

What does

```
f filename
```

do?

Printing the contents of the buffer — the Print command "p"

To *print* or list the contents of the buffer (or parts of it) on the terminal, we use the print command

```
p
```

The way this is done is as follows. We specify the lines where we want printing to begin and where we want it to end, separated by a comma, and followed by the letter "p". Thus to print the first two lines of the buffer, for example, (that is, lines 1 through 2) we say

```
1,2p (starting line=1, ending line=2 p)
```

Ed will respond with

```
Now is the time
for all good men
```

Suppose we want to print *all* the lines in the buffer. We could use "1,3p" as above if we knew there were exactly 3 lines in the buffer. But in general, we don't know how many there are, so what do we use for the ending line number? *Ed* provides a shorthand symbol for "line number of last line in buffer" — the dollar sign "\$". Use it this way:

```
1,$p
```

This will print *all* the lines in the buffer (line 1 to last line.) If you want to stop the printing before it is finished, push the DEL or Delete key; *ed* will type

```
?
```

and wait for the next command.

To print the *last* line of the buffer, we could use

```
,$p
```

but *ed* lets us abbreviate this to

\$p

We can print any single line by typing the line number followed by a "p". Thus

1p

produces the response

Now is the time

which is the first line of the buffer.

In fact, *ed* lets us abbreviate even further: we can print any single line by typing *just* the line number - no need to type the letter "p". So if we say

\$

ed will print the last line of the buffer for us.

We can also use "\$" in combinations like

\$-1,\$p

which prints the last two lines of the buffer. This helps when we want to see how far we got in typing.

Exercise 3:

As before, create some text using the append command and experiment with the "p" command. You will find, for example, that you can't print line 0 or a line beyond the end of the buffer, and that attempts to print a buffer in reverse order by saying

3,1p

don't work.

The current line - "Dot" or "."

Suppose our buffer still contains the six lines as above, that we have just typed

1,3p

and *ed* has printed the three lines for us. Try typing just

p (no line numbers).

This will print

to come to the aid of their party.

which is the third line of the buffer. In fact it is the last (most recent) line that we have done anything with. (We just printed it!) We can repeat this "p" command without line numbers, and it will continue to print line 3.

The reason is that *ed* maintains a record of the last line that we did anything to (in this case, line 3, which we just printed) so that it can be used instead of an explicit line number. This most recent line is referred to by the shorthand symbol

(pronounced "dot").

Dot is a line number in the same way that "\$" is; it means exactly "the current line", or loosely, "the line we most recently did something to." We can use it in several ways - one possibility is to say

.,\$p

This will print all the lines from (including) the current line to the end of the buffer. In our case these are lines 3 through 6.

Some commands change the value of dot, while others do not. The print command sets dot to the number of the last line printed; by our last command, we would have "." = "\$" = 6.

Dot is most useful when used in combinations like this one:

.,+1 (or equivalently, .+1p)

This means "print the next line" and gives us a handy way to step slowly through a buffer. We can also say

.,-1 (or .-1p)

which means "print the line *before* the current line." This enables us to go backwards if we wish. Another useful one is something like

.,-3,.-1p

which prints the previous three lines.

Don't forget that all of these change the value of dot. You can find out what dot is at any time by typing

.,=

Ed will respond by printing the value of dot.

Let's summarize some things about the "p" command and dot. Essentially "p" can be preceded by 0, 1, or 2 line numbers. If there is no line number given, it prints the "current line", the line that dot refers to. If there is one line number given (with or without the letter "p"), it prints that line (and dot is set there); and if there are two line numbers, it prints all the lines in that range (and sets dot to the last line printed.) If two line numbers are specified the first can't be bigger than the second (see Exercise 2.)

Typing a single newline will cause printing of the next line - it's equivalent to ".+1p". Try it. Try typing "^" - it's equivalent to ".-1p".

Deleting lines: the "d" command

Suppose we want to get rid of the three extra lines in the buffer. This is done by the *delete* command

d

Except that "d" deletes lines instead of printing them, its action is similar to that of "p". The lines to be deleted are specified for "d" exactly as they are for "p":

starting line, ending line d

Thus the command

4,\$d

deletes lines 4 through the end. There are now three lines left, as we can check by using

l,\$p

And notice that "\$" now is line 3! Dot is set to the next line after the last line deleted, unless the last line deleted is the last line in the buffer. In that case, dot is set to "\$".

Exercise 4:

Experiment with "a", "e", "r", "w", "p", and "d" until you are sure that you know what they do, and until you understand how dot, "\$", and line numbers are used.

If you are adventurous, try using line numbers with "a", "r", and "w" as well. You will find that "a" will append lines *after* the line number that you specify (rather than after dot); that "r" reads a file in *after* the line number you specify (not necessarily at the end of the buffer); and that "w" will write out exactly the lines you specify, not necessarily the whole buffer. These variations are sometimes handy. For instance you can insert a file at the beginning of a buffer by saying

Or filename

and you can enter lines at the beginning of the buffer by saying

0a
... text ...

Notice that ".w" is *very* different from

w

Modifying text: the Substitute command "s"

We are now ready to try one of the most important of all commands - the substitute command

s

This is the command that is used to change individual words or letters within a line or group of lines. It is what we use, for example, for correcting spelling mistakes and typing errors.

Suppose that by a typing error, line 1 says

Now is th time

- the "e" has been left off "the". We can use "s" to fix this up as follows:

1s/th/the/

This says: "in line 1, substitute for the characters 'th' the characters 'the'." To verify that it works (*ed* will not print the result automatically) we say

p

and get

Now is the time

which is what we wanted. Notice that dot must have been set to the line where the substitution took place, since the "p" command printed that line. Dot is always set this way with the "s" command.

The general way to use the substitute command is

starting-line, ending-line s/change this/to this/

Whatever string of characters is between the first pair of slashes is replaced by whatever is between the second pair, in *all* the lines between starting line and ending line. Only the first occurrence on each line is changed, however. If you want to change *every* occurrence, see Exercise 5. The rules for line numbers are the same as those for "p", except that dot is set to the last line changed. (But there is a trap for the unwary: if no substitution took place, dot is *not* changed. This causes an error "?" as a warning.)

Thus we can say

1,\$s/speling/spelling/

and correct the first spelling mistake on each line in the text. (This is useful for people who are consistent misspellers!)

If no line numbers are given, the "s" command assumes we mean "make the substitution on line dot", so it changes things only on the current line. This leads to the very common sequence

s/something/something else/p

which makes some correction on the current line, and then prints it, to make sure it worked out right. If it didn't, we can try again. (Notice that we put a print command on the same line as the substitute. With few exceptions, "p" can follow any command; no other multi-command lines are legal.)

It's also legal to say

s/...//

which means "change the first string of characters to *nothing*", i.e., remove them. This is useful for deleting extra words in a line or removing extra letters from words. For instance, if we had

Nowxx is the time

we can say

s/xx//p

to get

Now is the time

Notice that "/" here means "no characters", not a blank. There *is* a difference! (See below for another meaning of "/".)

Exercise 5:

Experiment with the substitute command. See what happens if you substitute for some word on a line with several occurrences of that word. For example, do this:

a
the other side of the coin
.
s/the/on the/p

You will get

on the other side of the coin

A substitute command changes only the first occurrence of the first string. You can change all occurrences by adding a "g" (for "global") to the "s" command, like this:

s/.../.../gp

Try other characters instead of slashes to delimit the two sets of characters in the "s" command - anything should work except blanks or tabs.

(If you get funny results using any of the characters

^ . \$ [* \

read the section on "Special Characters".)

Context searching - "/.../"

With the substitute command mastered, we can move on to another highly important idea of *ed* - context searching.

Suppose we have our original three line text in the buffer:

Now is the time
for all good men
to come to the aid of their party.

Suppose we want to find the line that contains "their" so we can change it to "the". Now with only three lines in the buffer, it's pretty easy to keep track of what line the word "their" is on. But if the buffer contained several hundred lines, and we'd been making changes, deleting and rearranging lines, and so on, we would no longer really know what this line number would be. Context searching is simply a method of specifying the desired line, regardless of what its number is, by specifying some context on it.

The way we say "search for a line that contains this particular string of characters" is to type

/string of characters we want to find/

For example, the *ed* line

/their/

is a context search which is sufficient to find the desired line - it will locate the next occurrence of the characters between slashes ("their"). It also sets dot to that line and prints the line for verification:

to come to the aid of their party.

"Next occurrence" means that *ed* starts looking for the string at line ".+1", searches to the end of the buffer, then continues at line 1 and searches to line dot. (That is, the search "wraps around" from "\$" to 1.) It scans all the lines in the buffer until it either finds the desired line or gets back to dot again. If the given string of characters can't be found in any line, *ed* types the error message

?

Otherwise it prints the line it found.

We can do both the search for the desired line *and* a substitution all at once, like this:

/their/s/their/the/p

which will yield

to come to the aid of the party.

There were three parts to that last command: context search for the desired line, make the substitution, print the line.

The expression `/their/` is a context search expression. In their simplest form, all context search expressions are like this — a string of characters surrounded by slashes. Context searches are interchangeable with line numbers, so they can be used by themselves to find and print a desired line, or as line numbers for some other command, like `s`. We used them both ways in the examples above.

Suppose the buffer contains the three familiar lines

```
Now is the time
for all good men
to come to the aid of their party.
```

Then the *ed* line numbers

```
/Now/+1
/good/
/party/-1
```

are all context search expressions, and they all refer to the same line (line 2). To make a change in line 2, we could say

```
/Now/+1s/good/bad/
```

or

```
/good/s/good/bad/
```

or

```
/party/-1s/good/bad/
```

The choice is dictated only by convenience. We could print all three lines by, for instance

```
/Now/,/party/p
```

or

```
/Now/,/Now/+2p
```

or by any number of similar combinations. The first one of these might be better if we don't know how many lines are involved. (Of course, if there were only three lines in the buffer, we'd use

```
1,$p
```

but not if there were several hundred.)

The basic rule is: a context search expression is *the same as* a line number, so it can be used wherever a line number is needed.

Exercise 6:

Experiment with context searching. Try a body of text with several occurrences of the same string of characters, and scan through it using the same context search.

Try using context searches as line numbers for the substitute, print and delete commands. (They can also be used with `r`, `w`, and `d`.)

Try context searching using `?text?` instead of `/text/`. This scans lines in the buffer in reverse order rather than normal. This is sometimes useful if you go too far while looking for some string of characters — it's an easy way to back up.

(If you get funny results with any of the characters

```
$ [ * \
```

read the section on "Special Characters".)

Ed provides a shorthand for repeating a context search for the same string. For example, the *ed* line number

```
/string/
```

will find the next occurrence of "string". It often happens that this is not the desired line, so the search must be repeated. This can be done by typing merely

```
//
```

This shorthand stands for "the most recently used context search expression." It can also be used as the first string of the substitute command, as in

```
/string1/s//string2/
```

which will find the next occurrence of "string1" and replace it by "string2". This can save a lot of typing. Similarly

```
??
```

means "scan backwards for the same expression."

Change and Insert — "c" and "i"

This section discusses the *change* command

```
c
```

which is used to change or replace a group of one or more lines, and the *insert* command

```
i
```

which is used for inserting a group of one or more lines.

"Change", written as

```
c
```

is used to replace a number of lines with different lines, which are typed in at the terminal. For example, to change lines `+.1` through `$$` to something else, type

```
+.1,$c
```

```
... type the lines of text you want here ...
```

The lines you type between the `c` command and the `.` will take the place of the original

lines between start line and end line. This is most useful in replacing a line or several lines which have errors in them.

If only one line is specified in the "c" command, then just that line is replaced. (You can type in as many replacement lines as you like.) Notice the use of "." to end the input - this works just like the "." in the append command and must appear by itself on a new line. If no line number is given, line dot is replaced. The value of dot is set to the last line you typed in.

"Insert" is similar to append - for instance

```
/string/i
... type the lines to be inserted here ...
```

will insert the given text *before* the next line that contains "string". The text between "i" and "." is *inserted before* the specified line. If no line number is specified dot is used. Dot is set to the last line inserted.

Exercise 7:

"Change" is rather like a combination of delete followed by insert. Experiment to verify that

```
start, end d
i
... text ...
```

is almost the same as

```
start, end c
... text ...
```

These are not *precisely* the same if line "\$" gets deleted. Check this out. What is dot?

Experiment with "a" and "i", to see that they are similar, but not the same. You will observe that

```
line-number a
... text ...
```

appends *after* the given line, while

```
line-number i
... text ...
```

inserts *before* it. Observe that if no line number is given, "i" inserts before line dot, while "a" appends after line dot.

Moving text around: the "m" command

The move command "m" is used for cutting and pasting - it lets you move a group of lines from one place to another in the buffer. Suppose we want to put the first three lines of the buffer at the end instead. We could do it by saying:

```
1,3w temp
$r temp
1,3d
```

(Do you see why?) but we can do it a lot easier with the "m" command:

```
1,3m$
```

The general case is

start line, end line m after this line

Notice that there is a third line to be specified - the place where the moved stuff gets put. Of course the lines to be moved can be specified by context searches; if we had

```
First paragraph
...
end of first paragraph.
Second paragraph
...
end of second paragraph.
```

we could reverse the two paragraphs like this:

```
/Second/,/second/m/First/-1
```

Notice the "-1" - the moved text goes *after* the line mentioned. Dot gets set to the last line moved.

The global commands "g" and "v"

The *global* command "g" is used to execute one or more *ed* commands on all those lines in the buffer that match some specified string. For example

```
g/peling/p
```

prints all lines that contain "peling". More usefully,

```
g/peling/s//pelling/gp
```

makes the substitution everywhere on the line, then prints each corrected line. Compare this to

```
1,$s/peling/pelling/gp
```

which only prints the last line substituted. Another subtle difference is that the "g" command does not give a "?" if "peling" is not found where the "s" command will.

There may be several commands (including "a", "c" "i" "r", "w", but not "g"); in that case, every line except the last must end with a backslash "\":

```
g/xxx/.-1s/abc/def\  
.+2s/ghi/jkl\  
.-2.p
```

makes changes in the lines before and after each line that contains "xxx", then prints all three lines.

The "v" command is the same as "g", except that the commands are executed on every line that does *not* match the string following "v":

```
v/ /d
```

deletes every line that does not contain a blank.

Special Characters

You may have noticed that things just don't work right when you used some characters like ".", "*", "\$", and others in context searches and the substitute command. The reason is rather complex, although the cure is simple. Basically, *ed* treats these characters as special, with special meanings. For instance, *in a context search or the first string of the substitute command only*,

```
/x.y/
```

means "a line with an x, any character, and a y," not just "a line with an x, a period, and a y." A complete list of the special characters that can cause trouble is the following:

```
^ . $ [ * \  
_
```

Warning: The backslash character \ is special to *ed*. For safety's sake, avoid it where possible. If you have to use one of the special characters in a substitute command, you can turn off its magic meaning temporarily by preceding it with the backslash. Thus

```
s/\.\.*/backslash dot star/
```

will change "\.*)" into "backslash dot star".

Here is a hurried synopsis of the other special characters. First, the circumflex "^" signifies the beginning of a line. Thus

```
/^string/
```

finds "string" only if it is at the beginning of a line: it will find

```
string
```

but not

```
the string...
```

The dollar-sign "\$" is just the opposite of the circumflex; it means the end of a line:

```
/string$/
```

will only find an occurrence of "string" that is at the end of some line. This implies, of course,

that

```
/^string$/
```

will find only a line that contains just "string", and

```
/^.$/
```

finds a line containing exactly one character.

The character ".", as we mentioned above, matches anything;

```
/x.y/
```

matches any of

```
x+y
```

```
x-y
```

```
x y
```

```
x.y
```

This is useful in conjunction with "*", which is a repetition character; "a*" is a shorthand for "any number of a's," so ".*" matches any number of anythings. This is used like this:

```
s/.*/stuff/
```

which changes an entire line, or

```
s/.*///
```

which deletes all characters in the line up to and including the last comma. (Since ".*" finds the longest possible match, this goes up to the last comma.)

"[" is used with "]" to form "character classes"; for example,

```
/[1234567890]/
```

matches any single digit — any one of the characters inside the braces will cause a match.

Finally, the "&" is another shorthand character - it is used only on the right-hand part of a substitute command where it means "whatever was matched on the left-hand side". It is used to save typing. Suppose the current line contained

```
Now is the time
```

and we wanted to put parentheses around it. We could just retype the line, but this is tedious. Or we could say

```
s/(/(/
```

```
s/)/)/
```

using our knowledge of "*" and "\$". But the easiest way uses the "&":

```
s/.*/(&)/
```

This says "match the whole line, and replace it by itself surrounded by parens." The "&" can be used several times in a line; consider using

```
s/.*/&? &!!/
```

to produce

```
Now is the time? Now is the time!!
```

We don't have to match the whole line, of course: if the buffer contains

```
the end of the world
```

we could type

```
/world/s//& is at hand/
```

to produce

```
the end of the world is at hand
```

Observe this expression carefully, for it illustrates how to take advantage of *ed* to save typing. The string `"/world/"` found the desired line; the shorthand `"/"` found the same word in the line; and the `"&"` saved us from typing it again.

The `"&"` is a special character only within the replacement text of a substitute command, and has no special meaning elsewhere. We can turn off the special meaning of `"&"` by preceding it with a `"\"`:

```
s/ampersand/\&/
```

will convert the word "ampersand" into the literal symbol "&" in the current line.

Summary of Commands and Line Numbers

The general form of *ed* commands is the command name, perhaps preceded by one or two line numbers, and, in the case of *e*, *r* and *w*, followed by a file name. Only one command is allowed per line, but a *p* command may follow any other command (except for *e*, *r*, *w* and *q*).

a (*append*) Add lines to the buffer (at line dot, unless a different line is specified). Appending continues until `"."` is typed on a new line. Dot is set to the last line appended.

c (*change*) Change the specified lines to the new text which follows. The new lines are terminated by a `"."`. If no lines are specified, replace line dot. Dot is set to last line changed.

d (*delete*) Delete the lines specified. If none are specified, delete line dot. Dot is set to the first undeleted line, unless `"$"` is deleted, in which case dot is set to `"$"`.

e (*edit*) Edit new file. Any previous contents of the buffer are thrown away, so issue a *w* beforehand if you want to save them.

f (*file*) Print remembered filename. If a name follows *f* the remembered name will be set to it.

g (*global*) *g/---/commands* will execute the com-

mands on those lines that contain `"---"`, which can be any context search expression.

i (*insert*) Insert lines before specified line (or dot) until a `"."` is typed on a new line. Dot is set to last line inserted.

m (*move*) Move lines specified to after the line named after *m*. Dot is set to the last line moved.

p (*print*) Print specified lines. If none specified, print line dot. A single line number is equivalent to "line-number *p*". A single newline prints `"+1"`, the next line.

q (*quit*) Exit from *ed*. Wipes out all text in buffer!!

r (*read*) Read a file into buffer (at end unless specified elsewhere.) Dot set to last line read.

s (*substitute*) *s/string1/string2/* will substitute the characters of 'string2' for 'string1' in specified lines. If no line is specified, make substitution in line dot. Dot is set to last line in which a substitution took place, which means that if no substitution took place, dot is not changed. *s* changes only the first occurrence of string1 on a line; to change all of them, type a `"g"` after the final slash.

v (*exclude*) *v/---/commands* executes "commands" on those lines that *do not* contain `"---"`.

w (*write*) Write out buffer onto a file. Dot is not changed.

`.=` (*dot value*) Print value of dot. (`"="` by itself prints the value of `"$"`.)


`!` (*temporary escape*)

Execute this line as a UNIX command.

`/----/` Context search. Search for next line which contains this string of characters. Print it. Dot is set to line where string found. Search starts at `"+1"`, wraps around from `"$"` to 1, and continues to dot, if necessary.

`?-----?` Context search in reverse direction. Start search at `"-1"`, scan to 1, wrap around to `"$"`.





Advanced Editing on UNIX

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

This paper is meant to help secretaries, typists and programmers to make effective use of the UNIX facilities for preparing and editing text. It provides explanations and examples of

- special characters, line addressing and global commands in the editor **ed**;
- commands for “cut and paste” operations on files and parts of files, including the **mv**, **cp**, **cat** and **rm** commands, and the **r**, **w**, **m** and **t** commands of the editor;
- editing scripts and editor-based programs like **grep** and **sed**.

Although the treatment is aimed at non-programmers, new UNIX users with any background should find helpful hints on how to get their jobs done more easily.



Advanced Editing on UNIX

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

1. INTRODUCTION

Although UNIX provides remarkably effective tools for text editing, that by itself is no guarantee that everyone will automatically make the most effective use of them. In particular, people who are not computer specialists — typists, secretaries, casual users — often use the system less effectively than they might.

This document is intended as a sequel to *A Tutorial Introduction to the UNIX Text Editor* [1], providing explanations and examples of how to edit with less effort. (You should also be familiar with the material in *UNIX For Beginners* [2].) Further information on all commands discussed here can be found in *The UNIX Programmer's Manual* [3].

Examples are based on observations of users and the difficulties they encounter. Topics covered include special characters in searches and substitute commands, line addressing, the global commands, and line moving and copying. There are also brief discussions of effective use of related tools, like those for file manipulation, and those based on **ed**, like **grep** and **sed**.

A word of caution. There is only one way to learn to use something, and that is to *use* it. Reading a description is no substitute for trying something. A paper like this one should give you ideas about what to try, but until you actually try something, you will not learn it.

2. SPECIAL CHARACTERS

The editor **ed** is the primary interface to the system for many people, so it is worthwhile to know how to get the most out of **ed** for the least effort.

The next few sections will discuss shortcuts and labor-saving devices. Not all of these will be instantly useful to any one person, of course, but a few will be, and the others should give you ideas to store away for future use. And as always, until you try these things, they will remain theoretical knowledge, not something you have confidence in.

The List command 'l'

ed provides two commands for printing the contents of the lines you're editing. Most people are familiar with **p**, in combinations like

```
l,$p
```

to print all the lines you're editing, or

```
s/abc/def/p
```

to change 'abc' to 'def' on the current line. Less familiar is the *list* command **l** (the letter 'l'), which gives slightly more information than **p**. In particular, **l** makes visible characters that are normally invisible, such as tabs and backspaces. If you list a line that contains some of these, **l** will print each tab as \Rightarrow and each backspace as \Leftarrow . This makes it much easier to correct the sort of typing mistake that inserts extra spaces adjacent to tabs, or inserts a backspace followed by a space.

The **l** command also 'folds' long lines for printing — any line that exceeds 72 characters is printed on multiple lines; each printed line except the last is terminated by a backslash \backslash , so you can tell it was folded. This is useful for printing long lines on short terminals.

Occasionally the **l** command will print in a line a string of numbers preceded by a backslash, such as $\backslash07$ or $\backslash16$. These combinations are used to make visible characters that normally don't print, like form feed or vertical tab or bell. Each such combination is a single character. When you see such characters, be wary — they may have surprising meanings when printed on some terminals. Often their presence means that your finger slipped while you were typing; you almost never want them.

The Substitute Command 's'

Most of the next few sections will be taken up with a discussion of the substitute command **s**. Since this is the command for changing the contents of individual lines, it probably has the most complexity of any **ed** command, and the most potential for effective use.

As the simplest place to begin, recall the meaning of a trailing **g** after a substitute command. With

```
s/this/that/
```

and

```
s/this/that/g
```

the first one replaces the *first* 'this' on the line with 'that'. If there is more than one 'this' on the line, the second form with the trailing **g** changes *all* of them.

Either form of the **s** command can be followed by **p** or **l** to 'print' or 'list' (as described in the previous section) the contents of the line:

```
s/this/that/p
s/this/that/l
s/this/that/gp
s/this/that/gl
```

are all legal, and mean slightly different things. Make sure you know what the differences are.

Of course, any **s** command can be preceded by one or two 'line numbers' to specify that the substitution is to take place on a group of lines. Thus

```
1,$s/mispell/misspell/
```

changes the *first* occurrence of 'mispell' to 'misspell' on every line of the file. But

```
1,$s/mispell/misspell/g
```

changes *every* occurrence in every line (and this is more likely to be what you wanted in this particular case).

You should also notice that if you add a **p** or **l** to the end of any of these substitute commands, only the last line that got changed will be printed, not all the lines. We will talk later about how to print all the lines that were modified.

The Undo Command 'u'

Occasionally you will make a substitution in a line, only to realize too late that it was a ghastly mistake. The 'undo' command **u** lets you 'undo' the last substitution: the last line that was substituted can be restored to its previous state by typing the command

```
u
```

The Metacharacter '.'

As you have undoubtedly noticed when you use **ed**, certain characters have unexpected meanings when they occur in the left side of a substitute command, or in a search for a particu-

lar line. In the next several sections, we will talk about these special characters, which are often called 'metacharacters'.

The first one is the period '.'. On the left side of a substitute command, or in a search with '/.../', '.' stands for *any* single character. Thus the search

```
/x.y/
```

finds any line where 'x' and 'y' occur separated by a single character, as in

```
x+y
x-y
x y
x.y
```

and so on. (We will use to stand for a space whenever we need to make it clear.)

Since '.' matches a single character, that gives you a way to deal with funny characters printed by **l**. Suppose you have a line that, when printed with the **l** command, appears as

```
.... th\07is ....
```

and you want to get rid of the `\07` (which represents the bell character, by the way).

The most obvious solution is to try

```
s/\07//
```

but this will fail. (Try it.) The brute force solution, which most people would now take, is to re-type the entire line. This is guaranteed, and is actually quite a reasonable tactic if the line in question isn't too big, but for a very long line, re-typing is a bore. This is where the metacharacter '.' comes in handy. Since '\07' really represents a single character, if we say

```
s/th.is/this/
```

the job is done. The '.' matches the mysterious character between the 'h' and the 'i', *whatever it is*.

Bear in mind that since '.' matches any single character, the command

```
s/././
```

converts the first character on a line into a '.', which very often is not what you intended.

As is true of many characters in **ed**, the '.' has several meanings, depending on its context. This line shows all three:

```
.s/././
```

The first '.' is a line number: the number of the line we are editing, which is called 'line dot'. (We will discuss line dot more in Section 3.) The

second '.' is a metacharacter that matches any single character on that line. The third '.' is the only one that really is an honest literal period. On the *right* side of a substitution, '.' is not special. If you apply this command to the line

Now is the time.

the result will be

.ow is the time.

which is probably not what you intended.

The Backslash '\'

Since a period means 'any character', the question naturally arises of what to do when you really want a period. For example, how do you convert the line

Now is the time.

into

Now is the time?

The backslash '\' does the job. A backslash turns off any special meaning that the next character might have; in particular, '\' converts the '.' from a 'match anything' into a period, so you can use it to replace the period in

Now is the time.

like this:

s/\./?/

The pair of characters '\.' is considered by **ed** to be a single real period.

The backslash can also be used when searching for lines that contain a special character. Suppose you are looking for a line that contains

.PP

The search

/.PP/

isn't adequate, for it will find a line like

THE APPLICATION OF ...

because the '.' matches the letter 'A'. But if you say

/\.PP/

you will find only lines that contain '.PP'.

The backslash can also be used to turn off special meanings for characters other than '.'. For example, consider finding a line that contains a backslash. The search

/\

won't work, because the '\' isn't a literal '\', but instead means that the second '/' no longer delimits the search. But by preceding a backslash with another one, you can search for a literal backslash. Thus

/\\

does work. Similarly, you can search for a forward slash '/' with

/\/

The backslash turns off the meaning of the immediately following '/' so that it doesn't terminate the */.../* construction prematurely.

As an exercise, before reading further, find two substitute commands each of which will convert the line

\x\y

into the line

\xy

Here are several solutions; verify that each works as advertised.

s/\\.\//
s/x../x/
s/..y/y/

A couple of miscellaneous notes about backslashes and special characters. First, you can use any character to delimit the pieces of an **s** command; there is nothing sacred about slashes. (But you must use slashes for context searching.) For instance, in a line that contains a lot of slashes already, like

//exec //sys.fort.go // etc...

you could use a colon as the delimiter — to delete all the slashes, type

s/::g

Second, if # and @ are your character erase and line kill characters, you have to type \# and \@; this is true whether you're talking to **ed** or any other program.

When you are adding text with **a** or **i** or **c**, backslash is not special, and you should only put in one backslash for each one you really want.

The Dollar Sign '\$'

The next metacharacter, the '\$', stands for 'the end of the line'. As its most obvious use, suppose you have the line

Now is the

and you wish to add the word 'time' to the end. Use the \$ like this:

s/\$/ time/

to get

Now is the time

Notice that a space is needed before 'time' in the substitute command, or you will get

Now is thetime

As another example, replace the second comma in this line with a period without altering the first.

Now is the time, for all good men,

The command needed is

s/,,\$/./

The \$ sign here provides context to make specific which comma we mean. Without it, of course, the s command would operate on the first comma to produce

Now is the time. for all good men,

As another example, to convert

Now is the time.

into

Now is the time?

as we did earlier, we can use

s/.\$/?/

Like '.', the '\$' has multiple meanings depending on context. In the line

\$s/\$/\$/

the first '\$' refers to the last line of the file, the second refers to the end of that line, and the third is a literal dollar sign, to be added to that line.

The Circumflex '^'

The circumflex (or hat or caret) '^' stands for the beginning of the line. For example, suppose you are looking for a line that begins with 'the'. If you simply say

/the/

you will in all likelihood find several lines that contain 'the' in the middle before arriving at the one you want. But with

/^the/

you narrow the context, and thus arrive at the desired one more easily.

The other use of '^' is of course to enable you to insert something at the beginning of a line:

s/^/ /

places a space at the beginning of the current line.

Metacharacters can be combined. To search for a line that contains *only* the characters

.PP

you can use the command

/^\.PP\$/

The Star '*'

Suppose you have a line that looks like this:

text x y text

where text stands for lots of text, and there are some indeterminate number of spaces between the x and the y. Suppose the job is to replace all the spaces between x and y by a single space. The line is too long to retype, and there are too many spaces to count. What now?

This is where the metacharacter '*' comes in handy. A character followed by a star stands for as many consecutive occurrences of that character as possible. To refer to all the spaces at once, say

s/x *y/x y/

The construction 'x*y' means 'as many spaces as possible'. Thus 'x*y' means 'an x, as many spaces as possible, then a y'.

The star can be used with any character, not just space. If the original example was instead

text x-----y text

then all '-' signs can be replaced by a single space with the command

s/x-*y/x y/

Finally, suppose that the line was

text x.....y text

Can you see what trap lies in wait for the unwary? If you blindly type

s/x.*y/x|y/

what will happen? The answer, naturally, is that it depends. If there are no other x's or y's on the line, then everything works, but it's blind luck, not good management. Remember that '.' matches any single character? Then '*' matches as many single characters as possible, and unless you're careful, it can eat up a lot more of the line than you expected. If the line was, for example, like this:

text x text x.....y text y text

then saying

s/x.*y/x|y/

will take everything from the first 'x' to the last 'y', which, in this example, is undoubtedly more than you wanted.

The solution, of course, is to turn off the special meaning of '.' with '\.':

s/x\.y/x|y/

Now everything works, for '\.*' means 'as many periods as possible'.

There are times when the pattern '*' is exactly what you want. For example, to change

Now is the time for all good men

into

Now is the time.

use '*' to eat up everything after the 'for':

s/ for.*./

There are a couple of additional pitfalls associated with '*' that you should be aware of. Most notable is the fact that 'as many as possible' means zero or more. The fact that zero is a legitimate possibility is sometimes rather surprising. For example, if our line contained

text xy text x y text

and we said

s/x|.*y/x|y/

the first 'xy' matches this pattern, for it consists of an 'x', zero spaces, and a 'y'. The result is that the substitute acts on the first 'xy', and does not touch the later one that actually contains some intervening spaces.

The way around this, if it matters, is to specify a pattern like

/x|.*y/

which says 'an x, a space, then as many more

spaces as possible, then a y'.

The other startling behavior of '*' is again related to the fact that zero is a legitimate number of occurrences of something followed by a star. The command

s/x*/y/g

when applied to the line

abcdef

produces

yaybycydyeyfy

which is almost certainly not what was intended. The reason for this behavior is that zero is a legal number of matches, and there are no x's at the beginning of the line (so that gets converted into a 'y'), nor between the 'a' and the 'b' (so that gets converted into a 'y'), nor ... and so on. Make sure you really want zero matches; if not, in this case write

s/xx*/y/g

'xx*' is one or more x's.

The Brackets '['']

Suppose that you want to delete any numbers that appear at the beginning of all lines of a file. You might first think of trying a series of commands like

1,\$s/^1*//

1,\$s/^2*//

1,\$s/^3*//

and so on, but this is clearly going to take forever if the numbers are at all long. Unless you want to repeat the commands over and over until finally all numbers are gone, you must get all the digits on one pass. This is the purpose of the brackets '[' and ']'. The construction

[1234567890]

matches any single digit — the whole thing is called a 'character class'. With a character class, the job is easy. The pattern '[0123456789]*' matches zero or more digits (an entire number), so

1,\$s/^[0123456789]*//

deletes all digits from the beginning of all lines.

Any characters can appear within a character class, and just to confuse the issue there are essentially no special characters inside the brackets; even the backslash doesn't have a special meaning. To search for special characters, for

example, you can say

```
/[\.\$~]/
```

Within [...], the '[' is not special. To get a '[' into a character class, make it the first character.

As a final frill on character classes, you can specify a class that means 'none of the following characters'. This is done by beginning the class with a '^':

```
[^1234567890]
```

stands for 'any character *except* a digit'. Thus you might find the first line that doesn't begin with a tab or space by a search like

```
/^[^ (space)(tab)]/
```

Within a character class, the circumflex has a special meaning only if it occurs at the beginning. Just to convince yourself, verify that

```
/^[^]/
```

finds a line that doesn't begin with a circumflex.

The Ampersand '&'

The ampersand '&' is used primarily to save typing. Suppose you have the line

```
Now is the time
```

and you want to make it

```
Now is the best time
```

Of course you can always say

```
s/the/the best/
```

but it seems silly to have to repeat the 'the'. The '&' is used to eliminate the repetition. On the *right* side of a substitute, the ampersand means 'whatever was just matched', so you can say

```
s/the/& best/
```

and the '&' will stand for 'the'. Of course this isn't much of a saving if the thing matched is just 'the', but if it is something truly long or awful, or if it is something like '.' which matches a lot of text, you can save some tedious typing. There is also much less chance of making a typing error in the replacement text. For example, to parenthesize a line, regardless of its length,

```
s/.*/(&)/
```

The ampersand can occur more than once on the right side:

```
s/the/& best and & worst/
```

makes

```
Now is the best and the worst time
```

and

```
s/.*/&? &!!/
```

converts the original line into

```
Now is the time? Now is the time!!
```

To get a literal ampersand, naturally the backslash is used to turn off the special meaning:

```
s/ampersand/\\&/
```

converts the word into the symbol. Notice that '&' is not special on the left side of a substitute, only on the *right* side.

Substituting Newlines

ed provides a facility for splitting a single line into two or more shorter lines by 'substituting in a newline'. As the simplest example, suppose a line has gotten unmanageably long because of editing (or merely because it was unwisely typed). If it looks like

```
text xy text
```

you can break it between the 'x' and the 'y' like this:

```
s/xy/x\  
y/
```

This is actually a single command, although it is typed on two lines. Bearing in mind that '\' turns off special meanings, it seems relatively intuitive that a '\' at the end of a line would make the newline there no longer special.

You can in fact make a single line into several lines with this same mechanism. As a large example, consider underlining the word 'very' in a long line by splitting 'very' onto a separate line, and preceding it by the **roff** or **nroff** formatting command '.ul'.

```
text a very big text
```

The command

```
s/()very()/\  
.ul\  
very\  
/
```

converts the line into four shorter lines, preceding the word 'very' by the line '.ul', and eliminating the spaces around the 'very', all at the same time.

When a newline is substituted in, dot is left pointing at the last line created.

Regrettably there is no way to go in the opposite direction: **ed** will not convert two lines into one.

Rearranging a Line with \ (... \)

(This section should be skipped on first reading.) Recall that '&' is a shorthand that stands for whatever was matched by the left side of an **s** command. In much the same way you can capture separate pieces of what was matched; the only difference is that you have to specify on the left side just what pieces you're interested in.

Suppose, for instance, that you have a file of lines that consist of names in the form

Smith, A. B.
Jones, C.

and so on, and you want the initials to precede the name, as in

A. B. Smith
C. Jones

It is possible to do this with a series of editing commands, but it is tedious and error-prone. (It is instructive to figure out how it is done, though.)

The alternative is to 'tag' the pieces of the pattern (in this case, the last name, and the initials), and then rearrange the pieces. On the left side of a substitution, if part of the pattern is enclosed between \ (and \), whatever matched that part is remembered, and available for use on the right side. On the right side, the symbol '\1' refers to whatever matched the first \(...\) pair, '\2' to the second \(...\) , and so on.

The command

```
1,$s/^([^\,]*),([^\,]*)/\2\1/
```

although hard to read, does the job. The first \(...\) matches the last name, which is any string up to the comma; this is referred to on the right side with '\1'. The second \(...\) is whatever follows the comma and any spaces, and is referred to as '\2'.

Of course, with any editing sequence this complicated, it's foolhardy to simply run it and hope. The global commands **g** and **v** discussed in section 4 provide a way for you to print exactly those lines which were affected by the substitute command, and thus verify that it did what you wanted in all cases.

3. LINE ADDRESSING IN THE EDITOR

The next general area we will discuss is that of line addressing in **ed**, that is, how you specify what lines are to be affected by editing commands. We have already used constructions like

```
1,$s/x/y/
```

to specify a change on all lines. And most users are long since familiar with using a single newline (or return) to print the next line, and with

```
/thing/
```

to find a line that contains 'thing'. Less familiar, surprisingly enough, is the use of

```
?thing?
```

to scan *backwards* for the previous occurrence of 'thing'. This is especially handy when you realize that the thing you want to operate on is back up the page from where you are currently editing.

The slash and question mark are the only characters you can use to delimit a context search, though you can use essentially any character in a substitute command.

Address Arithmetic

The next step is to combine the line numbers like '.', '\$', '/.../' and '?...?' with '+' and '-'. Thus

```
$-1
```

is a command to print the next to last line of the current file (that is, one line before line '\$'). For example, to recall how far you got in a previous editing session,

```
$-5,$p
```

prints the last six lines. (Be sure you understand why it's six, not five.) If there aren't six, of course, you'll get an error message.

As another example,

```
.-3,.+3p
```

prints from three lines before where you are now (at line dot) to three lines after, thus giving you a bit of context. By the way, the '+' can be omitted:

```
.-3,.3p
```

is absolutely identical in meaning.

Another area in which you can save typing effort in specifying lines is to use '-' and '+' as line numbers by themselves.

by itself is a command to move back up one line in the file. In fact, you can string several minus signs together to move back up that many lines:

moves up three lines, as does '-3'. Thus

-3,+3p

is also identical to the examples above.

Since '-' is shorter than '-1', constructions like

-.s/bad/good/

are useful. This changes 'bad' to 'good' on the previous line and on the current line.

'+' and '-' can be used in combination with searches using '/.../' and '?...?', and with '\$'. The search

/thing/--

finds the line containing 'thing', and positions you two lines before it.

Repeated Searches

Suppose you ask for the search

/horrible thing/

and when the line is printed you discover that it isn't the horrible thing that you wanted, so it is necessary to repeat the search again. You don't have to re-type the search, for the construction

//

is a shorthand for 'the previous thing that was searched for', whatever it was. This can be repeated as many times as necessary. You can also go backwards:

??

searches for the same thing, but in the reverse direction.

Not only can you repeat the search, but you can use '/' as the left side of a substitute command, to mean 'the most recent pattern'.

/horrible thing/

... ed prints line with 'horrible thing' ...

s//good/p

To go backwards and change a line, say

??s//good/

Of course, you can still use the '&' on the right hand side of a substitute to stand for whatever got matched:

//s//&&/p

finds the next occurrence of whatever you searched for last, replaces it by two copies of itself, then prints the line just to verify that it worked.

Default Line Numbers and the Value of Dot

One of the most effective ways to speed up your editing is always to know what lines will be affected by a command if you don't specify the lines it is to act on, and on what line you will be positioned (i.e., the value of dot) when a command finishes. If you can edit without specifying unnecessary line numbers, you can save a lot of typing.

As the most obvious example, if you issue a search command like

/thing/

you are left pointing at the next line that contains 'thing'. Then no address is required with commands like s to make a substitution on that line, or p to print it, or l to list it, or d to delete it, or a to append text after it, or c to change it, or i to insert text before it.

What happens if there was no 'thing'? Then you are left right where you were — dot is unchanged. This is also true if you were sitting on the only 'thing' when you issued the command. The same rules hold for searches that use '?...?'; the only difference is the direction in which you search.

The delete command d leaves dot pointing at the line that followed the last deleted line. When line '\$' gets deleted, however, dot points at the new line '\$'

The line-changing commands a, c and i by default all affect the current line — if you give no line number with them, a appends text after the current line, c changes the current line, and i inserts text before the current line.

a, c, and i behave identically in one respect — when you stop appending, changing or inserting, dot points at the last line entered. This is exactly what you want for typing and editing on the fly. For example, you can say

a

... text ...

... botch ...

(minor error)

.

s/botch/correct/

(fix botched line)

a

... more text ...

without specifying any line number for the substitute command or for the second append command. Or you can say

```

a
... text ...
... horrible botch ...      (major error)
.
c                          (replace entire line)
... fixed up line ...

```

You should experiment to determine what happens if you add *no* lines with **a**, **c** or **i**.

The **r** command will read a file into the text being edited, either at the end if you give no address, or after the specified line if you do. In either case, dot points at the last line read in. Remember that you can even say **Or** to read a file in at the beginning of the text. (You can also say **0a** or **1i** to start adding text at the beginning.)

The **w** command writes out the entire file. If you precede the command by one line number, that line is written, while if you precede it by two line numbers, that range of lines is written. The **w** command does *not* change dot: the current line remains the same, regardless of what lines are written. This is true even if you say something like

```
/^\.AB/,/^\.AE/w abstract
```

which involves a context search.

Since the **w** command is so easy to use, you should save what you are editing regularly as you go along just in case the system crashes, or in case you do something foolish, like clobbering what you're editing.

The least intuitive behavior, in a sense, is that of the **s** command. The rule is simple — you are left sitting on the last line that matched the pattern. If there were no matches, then dot is unchanged.

To illustrate, suppose that there are three lines in the buffer, and you are sitting on the middle one:

```

x1
x2
x3

```

Then the command

```
-,+s/x/y/p
```

prints the third line, which is the last one changed. But if the three lines had been

```

x1
y2
y3

```

and the same command had been issued while dot pointed at the second line, then the result would be to change and print only the first line, and that is where dot would be set.

Semicolon ';'

Searches with '/.../' and '?...?' start at the current line and move forward or backward respectively until they either find the pattern or get back to the current line. Sometimes this is not what is wanted. Suppose, for example, that the buffer contains lines like this:

```

.
.
.
ab
.
.
bc
.
.

```

Starting at line 1, one would expect that the command

```
/a/,b/p
```

prints all the lines from the 'ab' to the 'bc' inclusive. Actually this is not what happens. *Both* searches (for 'a' and for 'b') start from the same point, and thus they both find the line that contains 'ab'. The result is to print a single line. Worse, if there had been a line with a 'b' in it before the 'ab' line, then the print command would be in error, since the second line number would be less than the first, and it is illegal to try to print lines in reverse order.

This is because the comma separator for line numbers doesn't set dot as each address is processed; each search starts from the same place. In **ed**, the semicolon ';' can be used just like comma, with the single difference that use of a semicolon forces dot to be set at that point as the line numbers are being evaluated. In effect, the semicolon 'moves' dot. Thus in our example above, the command

```
/a;/b/p
```

prints the range of lines from 'ab' to 'bc', because after the 'a' is found, dot is set to that line, and then 'b' is searched for, starting beyond that line.

This property is most often useful in a very simple situation. Suppose you want to find the *second* occurrence of 'thing'. You could say

```
/thing/  
//
```

but this prints the first occurrence as well as the second, and is a nuisance when you know very well that it is only the second one you're interested in. The solution is to say

```
/thing://
```

This says to find the first occurrence of 'thing', set dot to that line, then find the second and print only that.

Closely related is searching for the second previous occurrence of something, as in

```
?something?:??
```

Printing the third or fourth or ... in either direction is left as an exercise.

Finally, bear in mind that if you want to find the first occurrence of something in a file, starting at an arbitrary place within the file, it is not sufficient to say

```
1:/thing/
```

because this fails if 'thing' occurs on line 1. But it is possible to say

```
0:/thing/
```

(one of the few places where 0 is a legal line number), for this starts the search at line 1.

Interrupting the Editor

As a final note on what dot gets set to, you should be aware that if you hit the interrupt or delete or rubout or break key while **ed** is doing a command, things are put back together again and your state is restored as much as possible to what it was before the command began. Naturally, some changes are irrevocable — if you are reading or writing a file or making substitutions or deleting lines, these will be stopped in some clean but unpredictable state in the middle (which is why it is not usually wise to stop them). Dot may or may not be changed.

Printing is more clear cut. Dot is not changed until the printing is done. Thus if you print until you see an interesting line, then hit delete, you are *not* sitting on that line or even near it. Dot is left where it was when the **p** command was started.

4. GLOBAL COMMANDS

The global commands **g** and **v** are used to perform one or more editing commands on all lines that either contain (**g**) or don't contain (**v**) a specified pattern.

As the simplest example, the command

```
g/UNIX/p
```

prints all lines that contain the word UNIX. The pattern that goes between the slashes can be anything that could be used in a line search or in a substitute command; exactly the same rules and limitations apply.

As another example, then,

```
g/\. /p
```

prints all the formatting commands in a file (lines that begin with '.').

The **v** command is identical to **g**, except that it operates on those line that do *not* contain an occurrence of the pattern. (There is no mnemonic significance to the letter 'v'.) So

```
v/\. /p
```

prints all the lines that don't begin with '.' — the actual text lines.

The command that follows **g** or **v** can be anything:

```
g/\. /d
```

deletes all lines that begin with '.', and

```
g/^$/d
```

deletes all empty lines.

Probably the most useful command that can follow a global is the substitute command, for this can be used to make a change and print each affected line for verification. For example, we could change the word UNIX to 'Unix' everywhere, and verify that it really worked, with

```
g/UNIX/s//Unix/gp
```

Notice that we used '/' in the substitute command to mean 'the previous pattern', in this case, UNIX. The **p** command is done on every line that matches the pattern, not just those on which a substitution took place.

The global command operates by making two passes over the file. On the first pass, all lines that match the pattern are marked. On the second pass, each marked line in turn is examined, dot is set to that line, and the command executed. This means that it is possible for the command that follows a **g** or **v** to use addresses, set dot, and so on quite freely.

```
g/\.PP/+
```

prints the line that follows each `.PP` command (the signal for a new paragraph in some formatting packages). Remember that `+` means 'one line past dot'. And

```
g/topic/?\SH?!
```

searches for each line that contains 'topic', scans backwards until it finds a line that begins `.SH` (a section heading) and prints the line that follows that, thus showing the section headings under which 'topic' is mentioned. Finally,

```
g/\.EQ/+ ,/\.EN/-p
```

prints all the lines that lie between `.EQ` and `.EN` formatting commands.

The `g` and `v` commands can also be preceded by line numbers, in which case the lines searched are only those in the range specified.

Multi-line Global Commands

It is possible to do more than one command under the control of a global command, although the syntax for expressing the operation is not especially natural or pleasant. As an example, suppose the task is to change 'x' to 'y' and 'a' to 'b' on all lines that contain 'thing'. Then

```
g/thing/s/x/y/\
s/a/b/
```

is sufficient. The `\` signals the `g` command that the set of commands continues on the next line; it terminates on the first line that does not end with `\`. (As a minor blemish, you can't use a substitute command to insert a newline within a `g` command.)

You should watch out for this problem: the command

```
g/x/s//y/\
s/a/b/
```

does *not* work as you expect. The remembered pattern is the last pattern that was actually executed, so sometimes it will be 'x' (as expected), and sometimes it will be 'a' (not expected). You must spell it out, like this:

```
g/x/s/x/y/\
s/a/b/
```

It is also possible to execute `a`, `c` and `i` commands under a global command; as with other multi-line constructions, all that is needed is to add a `\` at the end of each line except the last. Thus to add a `.nf` and `.sp` command

before each `.EQ` line, type

```
g/\.EQ/i\
.nf\
.sp
```

There is no need for a final line containing a `.` to terminate the `i` command, unless there are further commands being done under the global. On the other hand, it does no harm to put it in either.

5. CUT AND PASTE WITH UNIX COMMANDS

One editing area in which non-programmers seem not very confident is in what might be called 'cut and paste' operations — changing the name of a file, making a copy of a file somewhere else, moving a few lines from one place to another in a file, inserting one file in the middle of another, splitting a file into pieces, and splicing two or more files together.

Yet most of these operations are actually quite easy, if you keep your wits about you and go cautiously. The next several sections talk about cut and paste. We will begin with the UNIX commands for moving entire files around, then discuss `ed` commands for operating on pieces of files.

Changing the Name of a File

You have a file named 'memo' and you want it to be called 'paper' instead. How is it done?

The UNIX program that renames files is called `mv` (for 'move'); it 'moves' the file from one name to another, like this:

```
mv memo paper
```

That's all there is to it: `mv` from the old name to the new name.

```
mv oldname newname
```

Warning: if there is already a file around with the new name, its present contents will be silently clobbered by the information from the other file. The one exception is that you can't move a file to itself —

```
mv x x
```

is illegal.

Making a Copy of a File

Sometimes what you want is a copy of a file — an entirely fresh version. This might be because you want to work on a file, and yet save a copy in case something gets fouled up, or just

because you're paranoid.

In any case, the way to do it is with the `cp` command. (`cp` stands for 'copy'; UNIX is big on short command names, which are appreciated by heavy users, but sometimes a strain for novices.) Suppose you have a file called 'good' and you want to save a copy before you make some dramatic editing changes. Choose a name — 'savegood' might be acceptable — then type

```
cp good savegood
```

This copies 'good' onto 'savegood', and you now have two identical copies of the file 'good'. (If 'savegood' previously contained something, it gets overwritten.)

Now if you decide at some time that you want to get back to the original state of 'good', you can say

```
mv savegood good
```

(if you're not interested in 'savegood' anymore), or

```
cp savegood good
```

if you still want to retain a safe copy.

In summary, `mv` just renames a file; `cp` makes a duplicate copy. Both of them clobber the 'target' file if it already exists, so you had better be sure that's what you want to do *before* you do it.

Removing a File

If you decide you are really done with a file forever, you can remove it with the `rm` command:

```
rm savegood
```

throws away (irrevocably) the file called 'savegood'.

Putting Two or More Files Together

The next step is the familiar one of collecting two or more files into one big one. This will be needed, for example, when the author of a paper decides that several sections need to be combined into one. There are several ways to do it, of which the cleanest, once you get used to it, is a program called `cat`. (Not *all* UNIX programs have two-letter names.) `cat` is short for 'concatenate', which is exactly what we want to do.

Suppose the job is to combine the files 'file1' and 'file2' into a single file called 'bigfile'. If you say

```
cat file
```

the contents of 'file' will get printed on your terminal. If you say

```
cat file1 file2
```

the contents of 'file1' and then the contents of 'file2' will *both* be printed on your terminal, in that order. So `cat` combines the files, all right, but it's not much help to print them on the terminal — we want them in 'bigfile'.

Fortunately, there is a way. You can tell UNIX that instead of printing on your terminal, you want the same information put in a file. The way to do it is to add to the command line the character `>` and the name of the file where you want the output to go. Then you can say

```
cat file1 file2 >bigfile
```

and the job is done. (As with `cp` and `mv`, you're putting something into 'bigfile', and anything that was already there is destroyed.)

This ability to 'capture' the output of a program is one of the most useful aspects of UNIX. Fortunately it's not limited to the `cat` program — you can use it with *any* program that prints on your terminal. We'll see some more uses for it in a moment.

Naturally, you can combine several files, not just two:

```
cat file1 file2 file3 ... >bigfile
```

collects a whole bunch.

Question: is there any difference between

```
cp good savegood
```

and

```
cat good >savegood
```

Answer: for most purposes, no. You might reasonably ask why there are two programs in that case, since `cat` is obviously all you need. The answer is that `cp` will do some other things as well, which you can investigate for yourself by reading the manual. For now we'll stick to simple usages.

Adding Something to the End of a File

Sometimes you want to add one file to the end of another. We have enough building blocks now that you can do it; in fact before reading further it would be valuable if you figured out how. To be specific, how would you use `cp`, `mv` and/or `cat` to add the file 'good1' to the end of the file 'good'?

You could try

```
cat good good1 >temp
mv temp good
```

which is probably most direct. You should also understand why

```
cat good good1 >good
```

doesn't work. (Don't practice with a good 'good'!)

The easy way is to use a variant of >, called >>. In fact, >> is identical to > except that instead of clobbering the old file, it simply tacks stuff on at the end. Thus you could say

```
cat good1 >>good
```

and 'good1' is added to the end of 'good'. (And if 'good' didn't exist, this makes a copy of 'good1' called 'good'.)

6. CUT AND PASTE WITH THE EDITOR

Now we move on to manipulating pieces of files — individual lines or groups of lines. This is another area where new users seem unsure of themselves.

Filenames

The first step is to ensure that you know the **ed** commands for reading and writing files. Of course you can't go very far without knowing **r** and **w**. Equally useful, but less well known, is the 'edit' command **e**. Within **ed**, the command

```
e newfile
```

says 'I want to edit a new file called *newfile*, without leaving the editor.' The **e** command discards whatever you're currently working on and starts over on *newfile*. It's exactly the same as if you had quit with the **q** command, then re-entered **ed** with a new file name, except that if you have a pattern remembered, then a command like **//** will still work.

If you enter **ed** with the command

```
ed file
```

ed remembers the name of the file, and any subsequent **e**, **r** or **w** commands that don't contain a filename will refer to this remembered file. Thus

```
ed file1
... (editing) ...
w      (writes back in file1)
e file2 (edit new file, without leaving editor)
... (editing on file2) ...
w      (writes back on file2)
```

(and so on) does a series of edits on various files without ever leaving **ed** and without typing the

name of any file more than once.

You can find out the remembered file name at any time with the **f** command; just type **f** without a file name. You can also change the name of the remembered file name with **f**; a useful sequence is

```
ed precious
f junk
... (editing) ...
```

which gets a copy of a precious file, then uses **f** to guarantee that a careless **w** command won't clobber the original.

Inserting One File into Another

Suppose you have a file called 'memo', and you want the file called 'table' to be inserted just after the reference to Table 1. That is, in 'memo' somewhere is a line that says

```
Table 1 shows that ...
```

and the data contained in 'table' has to go there, probably so it will be formatted properly by **nroff** or **troff**. Now what?

This one is easy. Edit 'memo', find 'Table 1', and add the file 'table' right there:

```
ed memo
/Table 1/
Table 1 shows that ... [response from ed]
.r table
```

The critical line is the last one. As we said earlier, the **r** command reads a file; here you asked for it to be read in right after line dot. An **r** command without any address adds lines at the end, so it is the same as **\$r**.

Writing out Part of a File

The other side of the coin is writing out part of the document you're editing. For example, maybe you want to split out into a separate file that table from the previous example, so it can be formatted and tested separately. Suppose that in the file being edited we have

```
.TS
...[lots of stuff]
.TE
```

which is the way a table is set up for the **tbl** program. To isolate the table in a separate file called 'table', first find the start of the table (the '.TS' line), then write out the interesting part:

```
/^\.TS/
.TS [ed prints the line it found]
.,/^\.TE/w table
```

and the job is done. If you are confident, you can do it all at once with

```
/'\TS/;'\'TE/w table
```

The point is that the **w** command can write out a group of lines, instead of the whole file. In fact, you can write out a single line if you like; just give one line number instead of two. For example, if you have just typed a horribly complicated line and you know that it (or something like it) is going to be needed later, then save it — don't re-type it. In the editor, say

```
a
...lots of stuff...
...horrible line...
.
.w temp
a
...more stuff...
.
.r temp
a
...more stuff...
.
```

This last example is worth studying, to be sure you appreciate what's going on.

Moving Lines Around

Suppose you want to move a paragraph from its present position in a paper to the end. How would you do it? As a concrete example, suppose each paragraph in the paper begins with the formatting command **.PP**. Think about it and write down the details before reading on.

The brute force way (not necessarily bad) is to write the paragraph onto a temporary file, delete it from its current position, then read in the temporary file at the end. Assuming that you are sitting on the **.PP** command that begins the paragraph, this is the sequence of commands!

```
./'\.PP/-w temp
.//-d
$r temp
```

That is, from where you are now ('.') until one line before the next **.PP** (('/'\.PP/-') write onto 'temp'. Then delete the same lines. Finally, read 'temp' at the end.

As we said, that's the brute force way. The easier way (often) is to use the *move* command **m** that **ed** provides — it lets you do the whole set of operations at one crack, without any temporary file.

The **m** command is like many other **ed** commands in that it takes up to two line numbers in front that tell what lines are to be affected. It is also *followed* by a line number that tells where the lines are to go. Thus

```
line1, line2 m line3
```

says to move all the lines between 'line1' and 'line2' after 'line3'. Naturally, any of 'line1' etc., can be patterns between slashes, \$ signs, or other ways to specify lines.

Suppose again that you're sitting at the first line of the paragraph. Then you can say

```
./'\.PP/-m$
```

That's all.

As another example of a frequent operation, you can reverse the order of two adjacent lines by moving the first one to after the second. Suppose that you are positioned at the first. Then

```
m+
```

does it. It says to move line dot to after one line after line dot.

As you can see, the **m** command is more succinct and direct than writing, deleting and re-reading. When is brute force better anyway? This is a matter of personal taste — do what you have most confidence in. The main difficulty with the **m** command is that if you use patterns to specify both the lines you are moving and the target, you have to take care that you specify them properly, or you may well not move the lines you thought you did. The result of a botched **m** command can often be a mess. Doing the job a step at a time makes it easier for you to verify at each step that you accomplished what you wanted to. It's also a good idea to issue a **w** command before doing anything complicated; then if you goof, it's easy to back up to where you were.

Marks

ed provides a facility for marking a line with a particular name so you can later reference it by name regardless of its actual line number. This can be handy for moving lines, and for keeping track of them as they move. The *mark* command is **k**; the command

```
kx
```

marks the current line with the name 'x'. If a line number precedes the **k**, that line is marked. (The mark name must be a single lower case letter.) Now you can refer to the marked line

with the address

'x

Marks are most useful for moving things around. Find the first line of the block to be moved, and mark it with 'a'. Then find the last line and mark it with 'b'. Now position yourself at the place where the stuff is to go and say

'a,'bm.

Bear in mind that only one line can have a particular mark name associated with it at any given time.

Copying Lines

We mentioned earlier the idea of saving a line that was hard to type or used often, so as to cut down on typing time. Of course this could be more than one line; then the saving is presumably even greater.

ed provides another command, called **t** (for 'transfer') for making a copy of a group of one or more lines at any point. This is often easier than writing and reading.

The **t** command is identical to the **m** command, except that instead of moving lines it simply duplicates them at the place you named. Thus

1,\$t\$

duplicates the entire contents that you are editing. A more common use for **t** is for creating a series of lines that differ only slightly. For example, you can say

```

a
..... x ..... (long line)
.
t.          (make a copy)
s/x/y/     (change it a bit)
t.          (make third copy)
s/y/z/     (change it a bit)

```

and so on.

The Temporary Escape '!'

Sometimes it is convenient to be able to temporarily escape from the editor to do some other UNIX command, perhaps one of the file copy or move commands discussed in section 5, without leaving the editor. The 'escape' command **!** provides a way to do this.

If you say

!*any UNIX command*

your current editing state is suspended, and the

UNIX command you asked for is executed. When the command finishes, **ed** will signal you by printing another **!**; at that point you can resume editing.

You can really do *any* UNIX command, including another **ed**. (This is quite common, in fact.) In this case, you can even do another **!**.

7. SUPPORTING TOOLS

There are several tools and techniques that go along with the editor, all of which are relatively easy once you know how **ed** works, because they are all based on the editor. In this section we will give some fairly cursory examples of these tools, more to indicate their existence than to provide a complete tutorial. More information on each can be found in [3].

Grep

Sometimes you want to find all occurrences of some word or pattern in a set of files, to edit them or perhaps just to verify their presence or absence. It may be possible to edit each file separately and look for the pattern of interest, but if there are many files this can get very tedious, and if the files are big (more than three or four thousand lines) it is impossible because of limits in **ed**.

The program **grep** was invented to get around these limitations. The patterns that we have described in the paper are often called 'regular expressions', and 'grep' stands for

g/re/p

That describes exactly what **grep** does — it prints every line in a set of files that contains a particular pattern. Thus

grep 'thing' file1 file2 file3 ...

finds 'thing' wherever it occurs in any of the files 'file1', 'file2', etc. **grep** also indicates the file in which the line was found, so you can later edit it if you like.

The pattern represented by 'thing' can be any pattern you can use in the editor, since **grep** and **ed** use exactly the same mechanism for pattern searching. It is wisest always to enclose the pattern in the single quotes '...' if it contains any non-alphabetic characters, since many such characters also mean something special to the UNIX command interpreter (the 'shell'). If you don't quote them, the command interpreter will try to interpret them before **grep** gets a chance.

There is also a way to find lines that *don't* contain a pattern:

```
grep -v 'thing' file1 file2 ...
```

finds all lines that don't contain 'thing'. The -v must occur in the position shown. Given grep and grep -v, it is possible to do things like selecting all lines that contain some combination of patterns. For example, to get all lines that contain 'x' but not 'y':

```
grep x file... | grep -v y
```

(The notation | is a 'pipe', which causes the output of the first command to be used as input to the second command; see [2].)

Editing Scripts

If a fairly complicated set of editing operations is to be done on a whole set of files, the easiest thing to do is to make up a 'script', i.e., a file that contains the operations you want to perform, then apply this script to each file in turn.

For example, suppose you want to change every UNIX to Unix and every GCOS to Gcos in a large number of files. Then put into the file 'script' the lines

```
g/UNIX/s//Unix/g
g/GCOS/s//Gcos/g
w
q
```

Now you can say

```
ed file1 <script
ed file2 <script
...
```

This causes ed to take its commands from the prepared script. Notice that the whole job has to be planned in advance.

And of course by using the UNIX command interpreter, you can cycle through a set of files automatically, with varying degrees of ease.

Sed

sed ('stream editor') is a version of the editor with restricted capabilities but which is capable of processing unlimited amounts of input. Basically sed copies its input to its output, applying one or more editing commands to each line of input.

As an example, suppose that we want to do the UNIX to Unix part of the example given above, but without rewriting the files. Then the command

```
sed 's/UNIX/Unix/g' file1 file2 ...
```

applies the command 's/UNIX/Unix/g' to all lines from 'file1', 'file2', etc., and copies all lines

to the output. The advantage of using sed in such a case is that it can be used with input too large for ed to handle, and that all the output can be collected in one place, either in a file or perhaps piped into another program.

If the editing transformation is so complicated that more than one editing command is needed, commands can be supplied from a file, or on the command line, with a slightly more complex syntax. To take commands from a file, for example,

```
sed -f cmdfile input-files...
```

sed has further capabilities, including conditional testing and branching, which we cannot go into here.

Acknowledgement

I am grateful to Ted Dolotta for his careful reading and valuable suggestions.

References

- [1] Brian W. Kernighan, *A Tutorial Introduction to the UNIX Text Editor*, Bell Laboratories internal memorandum.
- [2] Brian W. Kernighan, *UNIX For Beginners*, Bell Laboratories internal memorandum.
- [3] Ken L. Thompson and Dennis M. Ritchie, *The UNIX Programmer's Manual*. Bell Laboratories, 1975.



PWB/UNIX
Shell Tutorial

J. R. Mashey

September 1977

PWB/UNIX Shell Tutorial

CONTENTS

1. INTRODUCTION	1
2. OVERVIEW OF THE UNIX ENVIRONMENT	1
2.1 File System 2	
2.2 Processes 2	
3. SHELL BASICS	3
3.1 Commands 3	
3.2 Redirection of Standard Input and Output 4	
3.3 Command Lines 4	
3.4 Generation of Argument Lists 6	
3.5 Quoting Mechanisms 6	
3.6 Examples 6	
3.7 How the Shell Finds Commands 7	
3.8 Changing the State of the Shell and the .profile File 7	
4. USING THE SHELL AS A COMMAND: SHELL PROCEDURES	8
4.1 Invoking the Shell 8	
4.2 Passing Arguments to the Shell 8	
4.3 Shell Variables 9	
4.4 Initialization of \$p and \$z by the .path File 10	
4.5 Control Structures 11	
4.6 Onintr: Interrupt Handling 14	
4.7 Special I/O Redirections 14	
4.8 Quoting Revisited 15	
4.9 Creation and Organization of Shell Procedures 15	
5. MISCELLANEOUS SUPPORTING COMMANDS	16
5.1 Echo: Simple Output 16	
5.2 Pump: Shell Data Transfer 16	
5.3 Expr: Expression Evaluation 17	
5.4 Logname, Logdir, Logtty: Login Data 17	
6. EXAMPLES OF SHELL PROCEDURES	18
7. EFFECTIVE AND EFFICIENT SHELL PROGRAMMING	23
7.1 Overall Approach 23	
7.2 Approximate Measures of Resource Consumption 23	
7.3 Efficient Organization 24	
ACKNOWLEDGMENTS	25
REFERENCES	25

PWB/UNIX Shell Tutorial

J. R. Mashey

Bell Laboratories
Murray Hill, New Jersey 07974

The command language for PWB/UNIX* is a high-level programming language that is an extended version of the UNIX Shell. By utilizing the Shell as a programming language, one can eliminate much of the programming drudgery that often accompanies a large project. Many manual procedures can be quickly, cheaply, and conveniently automated. Because it is so easy to create and use Shell procedures, individual users and entire projects can customize the general PWB/UNIX environment into one tailored to their own respective requirements, organizational structure, and terminology.

This paper is actually a combination of several tutorials, as explained in {1}.¹ Some sections provide a basic tutorial for relatively new users. Other sections are intended for more experienced users and introduce them to Shell programming. Finally, some hints on programming techniques and efficiency are offered for those who make especially heavy use of Shell programming.

The accuracy of this tutorial is guaranteed *only* for the Shell of PWB/UNIX—Edition 1.0. Other versions of UNIX have other Shells. Although many of the basic concepts are similar, there exist many differences in features, especially those used to support Shell programming.

1. INTRODUCTION

In any programming project, some effort is used to build the end product. The remainder is consumed in building the supporting tools and procedures used to manage and maintain the end product. The second effort can far exceed the first, especially in larger projects. A good command language can be an invaluable tool for such projects. If it is a flexible programming language, it can be used to solve many internal support problems, without requiring compilable programs to be written, debugged, and maintained; its most important advantage is the ability to get the job done *now*. For a perspective on the motivations for using a command language in this way, see [1,2,6].

When users log into a PWB/UNIX system, they communicate with an instance of the Shell that reads commands typed at the terminal and arranges for their execution. Thus, the Shell's most important function is to provide a good interface for human beings. In addition, a sequence of commands may be preserved for later use by saving them in a file, called a *Shell procedure*, a *command file*, or a *runcom*, according to local preferences.

Some users need little knowledge of the Shell to do their work; others choose to make heavy use of its programming features. This tutorial may be read in several different ways, depending on the reader's interests. A brief discussion of the PWB/UNIX environment is found in {2}. The discussion in {3} covers aspects of the Shell that are important for everyone, while all of {4} and most of {5} are mainly of interest to those who write Shell procedures. A group of annotated Shell procedures is given in {6}. Finally, a brief discussion of efficiency is offered in {7}. This is found in its proper place (the end), and is intended for those who write especially time-consuming Shell procedures.

Complete beginners should *not* be reading this tutorial, but should work their way through other available tutorials first. See [7] for an appropriate plan of study. All the *commands* mentioned below are described in Section I of the *PWB/UNIX User's Manual* [3], while *system calls* are described in Section II and *subroutines* in Section III thereof.

2. OVERVIEW OF THE UNIX ENVIRONMENT

Full understanding of some later discussions depends on familiarity with PWB/UNIX; [9] is most useful for that, and it would be helpful to read at least one of [4,5,10]. For completeness, a short overview of the most relevant concepts is given below.

* UNIX is a Trademark/Service Mark of the Bell System.

1. The notation {*n*} refers to Section *n* of this tutorial.

2.1 File System

The PWB/UNIX file system's overall structure is that of a rooted tree composed of *directories* and other files. A *file name* is a sequence of characters. A *pathname* is a sequence of directory names followed by a file name, each separated from the previous one by a slash (/). If a pathname begins with a "/", the search for the file begins at the root of the entire tree; otherwise, it begins at the user's *current directory* (also known as the *working directory*). (The first kind of name is often called a *full pathname* because it is invariant with regard to the user's current directory.) The user may change the current directory at any time by using the *cd* or *chdir* command. In most cases, a file name and its corresponding pathname may be used interchangeably. Some sample names are:

- . name of the current directory.
- .. name of the parent directory of the current directory.
- / root directory of the entire file structure.
- /bin directory containing most of the frequently-used public commands.
- /a1/tf/jtb/bin a full pathname typical of multi-person programming projects. This one happens to be a private directory of commands belonging to person "jtb" in project "tf"; "a1" is the name of a *file system*.
- bin/umail a name depending on the current directory: it names file "umail" in subdirectory "bin" of the current directory. If the current directory is "/", it names "/bin/umail". If the current directory is "/a1/tf/jtb", it names "/a1/tf/jtb/bin/umail".
- memox name of a file in the current directory.

2.2 Processes

■ *Beginners should skip this section on first reading.*

An *image* is a computer execution environment, including memory image, register values, current directory, status of open files, information recorded at login time, and various other items. A *process* is the execution of an image; most PWB/UNIX commands execute as separate processes. One process may spawn another using the *fork* system call, which duplicates the image of the original (*parent*) process. The new (*child*) process continues to execute the same program as the parent. The two images are identical, except that the program can determine whether it is executing as parent or child. The program may continue execution of the image or may abandon it by issuing an *exec* system call, thus initiating execution of another program. In any case, each process is free to proceed in parallel with the other, although the parent quite commonly issues a *wait* system call to suspend execution until a child exits.

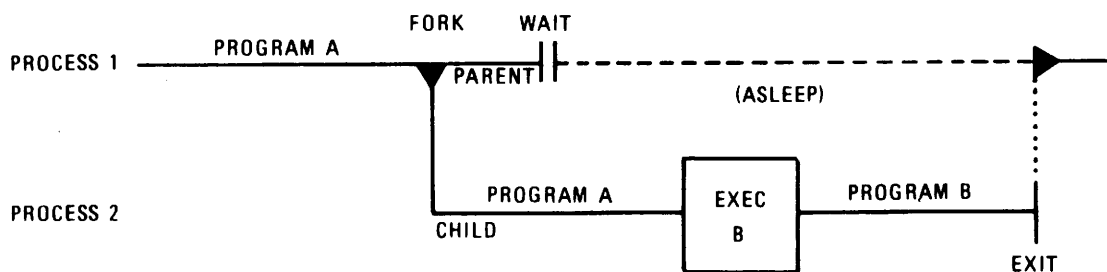


Figure 1

Figure 1 illustrates these ideas. *Program A* is executing (as *process 1*) and wishes to run *program B*. It "forks" and spawns a child (*process 2*) that continues to execute *program A*. The child abandons *A* by executing *B*, while the parent goes to sleep until the child exits.

A child inherits its parent's *open files*. This mechanism permits processes to share a common input stream in various ways. In particular, an open file possesses a *pointer* that indicates a position in the file and is modified by various operations. *Read* and *write* system calls copy a requested number of bytes from or to a file, beginning at the position given by the current value of the pointer. As a side effect, the pointer is incremented by the number of bytes transferred, yielding the effect of sequential I/O. *Seek* can be used to obtain random-access I/O; it sets the pointer to an absolute position within the file, or to a position offset either from the end of the file or from the current pointer position.

When a process terminates, it can set an eight-bit *return code* (or *exit code*) that is available to its parent. This code is usually used to indicate success or failure.

Signals indicate the occurrence of events that may have some impact on a process. A signal may be sent to a process by another process, from the keyboard, or by PWB/UNIX itself. For most types of signals, a process can arrange to be terminated on receipt of a signal, to ignore it completely, or to "catch" it and take appropriate action [4.6]. For example, an *interrupt* signal may be sent by depressing an appropriate key ("del", "break", or "rubout"). The action taken depends on the requirements of the specific program being executed:

- The Shell invokes most commands in such a way that they immediately die when an interrupt is received. For example, *pr* normally dies, allowing the user to terminate unwanted output.
- The Shell itself ignores interrupts when reading from the terminal, because it should continue execution even when the user terminates a command like *pr*.
- The editor *ed* chooses to "catch" interrupts so that it can halt its current action (especially printing) without terminating completely.

Limiting interprocess communication to a small number of well-defined methods is a great aid to uniformity, understandability, and reliability of programs. It encourages the "packaging" of each function into a small program that is easily connected to other programs, but depends very little on the internal workings of other programs.

3. SHELL BASICS

The *Shell* (i.e., the *sh* command) implements the command language visible to most PWB/UNIX users. It reads input from a terminal or a file and arranges for the execution of the requested commands. It is a small program (about forty pages of C code); many of its functions are actually provided by independent programs that work with it. It is *not* part of the operating system, but is an ordinary user program. The discussion below is adapted from [10,11].

3.1 Commands

A *command* is a sequence of non-blank arguments separated by blanks or tabs. The first argument (numbered *zero*) specifies the name of the command to be executed; any remaining arguments are passed as character-strings to the command executed. A command may be as simple as:

```
who
```

which prints the login names of logged-in users. The following line requests the *pr* command to print files a, b, and c:

```
pr a b c
```

If the first argument names a file that is *executable*² and is actually a load module, the Shell (as parent) spawns a new (child) process that immediately executes that program. If the file is marked executable, but is neither a load module nor a directory, it is assumed to be a *Shell procedure*, i.e., a file of ordinary text containing Shell command lines and possibly lines to be read by other programs. In this case, the Shell spawns a new instance of itself to read the file and execute the commands included in it. The following command requests that the on-line *PWB/UNIX User's Manual* [3] pages for the *who* and *pr* commands be printed on the terminal (the *man* command is actually implemented as a Shell procedure):

```
man who pr
```

2 As evidenced by an appropriate set of permission bits associated with that file.

From the user's viewpoint, executable programs and Shell procedures are invoked in exactly the same way. The Shell determines which implementation has been used, rather than requiring the user to do so. This preserves the uniformity of invocation and the ease of changing the implementation choice for a given command. The actions of the Shell in executing any of these commands are illustrated in Figure 1 [2.2].

3.2 Redirection of Standard Input and Output

When a command begins execution, it usually expects that three files are already open, a "standard input", a "standard output", and a "diagnostic output". When the user's original Shell is started, all three have already been opened to the user's terminal. A child process normally inherits these files from its parent. The Shell permits them to be redirected elsewhere before control is passed to an invoked command.

An argument to the Shell of the form "<file" or ">file" opens the specified file as standard input or output, respectively. An argument of the form ">>file" opens the standard output to the end of the file, thus providing a way to append data to it. In either output case, the Shell creates the file if it did not already exist. The following appends to file "log" the list of users who are logged in:

```
who >>log
```

In general, most commands neither know nor care whether their input (output) is coming from (going to) a terminal or file. Thus, commands can be used conveniently in many different contexts. A few commands vary their actions depending on the nature of their input or output, either for efficiency's sake, or to avoid useless actions (such as attempting random-access I/O on a terminal).

Redirection of the diagnostic output is discussed in [4.7.3].

3.3 Command Lines

A sequence of commands separated by "|" (or "^") make up a *pipeline*. Each command is run as a separate process connected to its neighbor(s) by *pipes*, i.e., the output of each command (except the last one) becomes the input of the next command in line. A *filter* is a command that reads its input, transforms it in some way, then writes it as output. A pipeline normally consists of a series of filters. Although the processes in a pipeline are permitted to execute in parallel, they are synchronized to the extent that each program needs to read the output of its predecessor. Many commands operate on individual lines of text, reading a line, processing it, writing it, and looping back for more input. Some must read larger amounts of data before producing output; *sort* is an example of the extreme case that requires all input to be read before any output is produced.

The following is an example of a typical pipeline: *nroff* is a text formatter whose output may contain reverse line motions; *col* converts these motions to a form that can be printed on a terminal lacking reverse motion capability; *reform* is used here to speed printing by converting the (tab-less) output of *col* to an equivalent one containing horizontal tab characters. The flag "-mm" indicates one of the more-commonly used formatting options, and "text" is the name of the file to be formatted:

```
nroff -mm text | col | reform
```

Figure 2 shows the sequence of actions that set up this pipeline. Not shown are actions by the Shell that create pipes and manipulate open files, causing the commands to be tied together correctly.

A *command line* consists of zero or more pipelines separated by semicolons or ampersands. If the last command in a pipeline is terminated by a semicolon (;) or a new-line character, the Shell waits for the command to finish before continuing to read command lines. It does *not* wait if the pipeline is terminated by an ampersand (&); both sequential and asynchronous execution are thus allowed. An asynchronous pipeline continues execution until it terminates voluntarily, or until its processes are *killed*. The first example below executes *who*, waits for it to terminate, and then executes *date*; the second invokes both commands in order, but does not wait for either one to finish. Figure 3 shows the actions of the Shell involved in executing these examples:

```
who >log; date  
who >log& date&
```

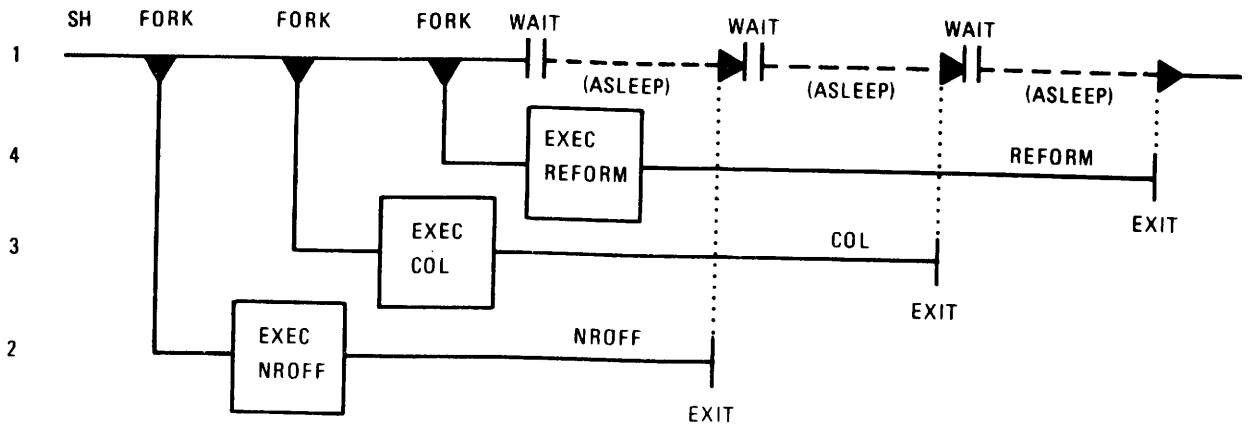


Figure 2

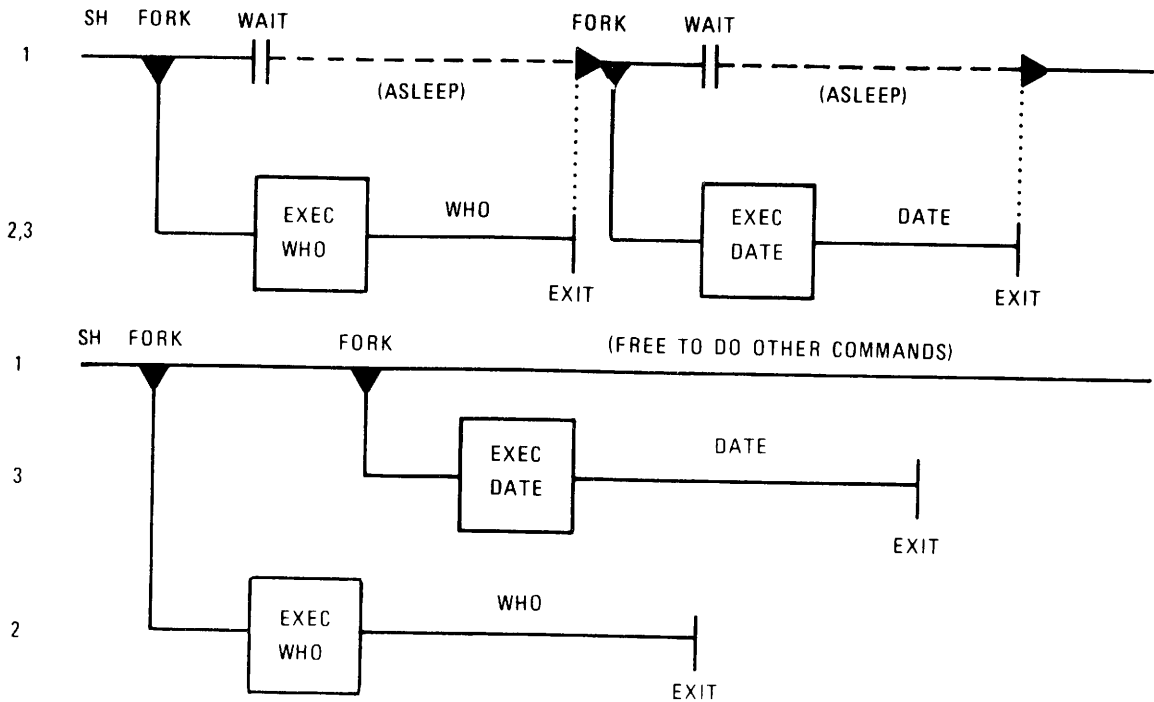


Figure 3

More typical uses of "&" include off-line printing, background compilation, and generation of jobs to be sent to other computers. For example:

```
nohup cc prog.c&
```

You continue working while the C compiler runs in background.

A command terminated by "&" is immune to interrupts, but it is wise to make it immune to hang-ups as well. The *nohup* command is used for this purpose. Without *nohup*, if you hang up while *cc* (the C compiler) is still executing, *cc* will be killed and your output will disappear.

➤ The "&" operator should be used with restraint, especially on heavily-loaded systems. Other users will not consider you a good citizen if you start up a large number of simultaneous, asynchronous processes without a compelling reason for doing so.

A simple command in a pipeline may be replaced by a command line enclosed in parentheses "()"; in this case, another instance of the Shell is spawned to execute the commands so enclosed. This action is

helpful in combining the output of several sequentially executed commands into a stream to be processed by a pipeline. The following line prints two separate documents in a way similar to that shown in a previous example:

```
(nroff -mm text1; nroff -mm text2) | col | reform
```

3.4 Generation of Argument Lists

Many command arguments are names of files. When certain characters are found in an argument, they cause replacement of that argument by a sorted list of zero or more file names obtained by pattern-matching on the contents of a directory. Most characters match themselves. The “?” matches *any one character*; the “*” matches *any string* of any characters (other than “/”), including the null string. Enclosing a set of characters within square brackets “[...]” causes the construct to match *any one character* in that set.³ Inside square brackets, a pair of characters separated by “-” includes in the set all characters lexically within the inclusive range of that pair.

For example, “*” matches all files in the current directory, “*tmp*” matches all names containing “tmp”, “[a-f]*” matches all files whose names begin with “a” through “f”, “*.c” matches all files ending in “.c”, and “/a1/tf/bin/?” matches all single-character names found in “/a1/tf/bin”. This capability saves much typing, and more importantly, makes it possible to organize information as large collections of small files that are named in disciplined ways.

Pattern-matching has several restrictions. If the first character of a file name is “.”, it can be matched only by an argument that begins with “.”. Pattern-matching is currently restricted to the last component in a pathname—the string “/a1/tf/*” is legal, but the string “/a1/*bin” is not. Pattern-matching does not apply to the name of the invoked command (i.e., argument number 0).

3.5 Quoting Mechanisms

If a character has a special meaning to the Shell, that meaning may be removed by preceding the character with a back-slash (\), the “\” acting as an escape and disappearing. A “\” followed by a new-line character is treated as a blank, permitting continuation of commands on additional input lines. A sequence of characters enclosed in single quotes (‘...’) is taken literally—“what you see is what you get”. The beginner should use single quotes in most instances. Double quotes (“...”) are required in a few cases, primarily inside Shell procedures. Double quotes hide the significance of most special characters, but allow substitution of Shell arguments and variables; see [4.8] for further details.

3.6 Examples

The following examples illustrate the variety of effects that can be obtained by combining a few commands in the ways described above. It may be helpful to try these examples at a terminal:

- who
Print (on the terminal) the list of logged-in users.
- who >>log
Append the list of logged-in users to the end of file “log”.
- who | wc -l
Print the number of logged-in users. (The argument to wc is “minus ell”.)
- who | pr
Print a paginated list of logged-in users.
- who | sort
Print an alphabetized list of logged-in users.
- who | grep pw
Print the list of logged-in users whose login name contains “pw”.
- who | grep pw | sort
Print an alphabetized list of logged-in users whose login names contain “pw”.

3. Be warned that square brackets are also mentioned in an entirely different section. In descriptions of commands, they indicate that the enclosed argument is optional.

- (date; who | wc -l) >>log
Append (to “log”) the current date followed by the count of logged-in users.
- who | sed 's/ .*//' | sort | uniq -d
Print only the login names of all users who are logged in more than once.

The *who* command does not *by itself* provide options to yield all these results—they are obtained by combining it with other commands. The kinds of operations illustrated above may be used in other circumstances; *who* just serves as the data source in these examples. As an exercise, replace “who |” by “</etc/passwd” in the above examples to see how a file can be used as a data source in the same way.

3.7 How the Shell Finds Commands

The Shell normally searches for commands in a way that permits them to be found in three distinct locations in the file structure. The Shell first attempts to use the command name as given; if this fails, it prepends the string “/bin/” to the name, and, finally, “/usr/bin/”. The effect is to search, in order, the current directory, “/bin”, and “/usr/bin”. For example, the *pr* and *man* commands are actually located in files “/bin/pr” and “/usr/bin/man”, respectively. A more complex pathname may be given, either to locate a file relative to the user’s current directory, or to access a command via an absolute pathname. If a command name *as given* contains a “/” (e.g., “/bin/sort” or “../cmd”), the prepending is *not* performed. Instead, a single attempt is made to execute the unmodified command name.

This mechanism gives the user a convenient way to execute public commands and commands in or “near” the current directory, as well as the ability to execute *any* accessible command regardless of its location in the file structure. Because the current directory is usually searched first, anyone can possess a private version of a public command without interfering with other users. Similarly, the creation of a new public command will not affect a user who already has a private command with the same name. This mechanism may be overridden {4.4}.

3.8 Changing the State of the Shell and the .profile File

The state of a given instance of the Shell may be altered in various ways. The following commands are used more often at the terminal than in Shell procedures.

The *cd* command (or its synonym *chdir*) changes the current directory of the Shell to the one specified. This can (and should) be used to change to a convenient place in the directory structure; *cd* is often combined with “()” to cause a sub-Shell to change to a different directory and execute some commands, without affecting the original Shell. The first sequence below extracts the component files of the archive file “/a1/tf/q.a” and places them in whatever directory is the current one; the second places them in directory “/a1/tf”:

```
ar x /a1/tf/q.a
(cd /a1/tf; ar x q.a)
```

The *opt* command sets various flags in the Shell. For example, “opt -p prompt-str” changes the Shell’s interactive prompt sequence from “% ” to *prompt-str*.⁴ Typing “opt -v” causes the Shell to enter verbose mode, in which it prints each command line before executing it {4.1}. Try this at the terminal to see how the Shell scans arguments. The output can be turned off by typing “opt +v”.

The *login* command causes the Shell to execute the *login* program directly, permitting a new login without re-dialing. A related command is *su*, which permits you to act with someone else’s access permissions *without* making you login again.

Wait causes the Shell to suspend execution until all of its child processes have terminated. It is used to assure termination of asynchronous processes.

When you login or use *su*, the Shell is invoked to read your commands, but if your current directory contains a file named “.profile”, the Shell reads it *before* reading commands from your terminal; “.profile” often contains commands that set tab stops and terminal delays, read mail, etc. See “.profile” in {6}.

4. The default prompt string “% ” is inconvenient for certain display (CRT) terminals.

4. USING THE SHELL AS A COMMAND: SHELL PROCEDURES

4.1 Invoking the Shell

The Shell is an ordinary command and may be invoked in the same way as other commands:

- sh *file* [*args*] A new instance of the Shell is explicitly invoked to read *file*. Arguments, if any, can be manipulated as described in {4.2}.
- sh -v *file* [*args*] This is equivalent to putting “opt -v” at the beginning of *file*. Each command line in *file* is printed before it is executed, thus tracing the progress of execution. *This is an important debugging aid.*
- file* [*args*] If *file* is marked executable, and is neither a directory nor a load module, the effect is that of “sh *file* [*args*]”, except that *file* may be found by the search procedure described in {3.7}.

4.2 Passing Arguments to the Shell

When a command line is scanned, any character sequence of the form $\$n$ is replaced by the n th argument to the Shell, counting the name of the file being read as $\$0$. This notation permits direct reference to the file name and up to 9 arguments. Additional arguments can be processed using the *shift* command. It shifts arguments to the left; i.e., the value of $\$1$ is thrown away, $\$2$ replaces $\$1$, $\$3$ replaces $\$2$, etc.; the rightmost argument becomes null. For example, consider the (executable) file “ripple” below. *Echo* writes its arguments to the standard output; *if*, *exit*, and *goto* are discussed later, but perform fairly obvious functions.⁵ The form “ $\$1$ ” is used rather than ‘ $\$1$ ’ because it is the value of the first argument that is desired, rather than the literal two-character string “ $\$1$ ”:

```
: loop
if "$1" = "" exit
echo $1 $2 $3 $4 $5 $6 $7 $8 $9
shift
goto loop
```

If the file were invoked by “ripple a b c”, it would print:

```
a b c
b c
c
```

The “shift n ” form of *shift* has no effect on the arguments to the left of the n th argument; the n th argument is discarded, and the higher-numbered ones shifted. Thus, “shift” is equivalent to “shift 1” (as is “shift 0”).

The notation $\$*$ causes substitution of *all* current arguments except $\$0$. Thus, the *echo* line in the “ripple” example above could be written in a better way as:

```
echo $*
```

These two *echo* commands are *not* equivalent: the first prints at most nine arguments; the second prints all its arguments. The $\$*$ notation is more concise and is less error-prone. One obvious application is in passing an arbitrary number of arguments to the *nroff* text formatter:

```
nroff -h -rT1 -T450 -mm $*
```

It is important to understand the sequence of actions used by the Shell in substituting arguments. First, the Shell reads one line of input, making all substitutions in a single pass; no rescanning is performed. Second, the Shell parses the resulting line. Third, the Shell executes all of the commands in that line. Thus, it is impossible for a command in a line to affect the argument values substituted into that same line. For example, the following sequence prints the same value twice, because the *shift* has no effect on the line in which it appears:

```
echo $1; shift; echo $1
```

⁵ Much better ways of coding this procedure are shown later. Lines that begin with “:” are labels and/or comments {4.5.1}.

On the other hand, the next sequence prints the first argument, followed by the second:

```
echo $1
shift
echo $1
```

4.3 Shell Variables

The Shell provides 26 string variables, \$a through \$z. Those in the range \$a through \$m are initialized to null strings at the beginning of execution and are never modified except by explicit user request. Some variables in the range \$n through \$z have specific initial values and may possibly be changed implicitly by the Shell during execution. A variable is assigned a value as follows:

```
= letter [ arg1 [ arg2 ] ]
```

If *arg1* is given, its value is assigned to the variable corresponding to *letter*. If two arguments are given, and if *arg1* is a null string, the value of *arg2* is assigned to the variable, permitting a convenient default mechanism. If neither *arg1* nor *arg2* are given, a single line is read from the standard input, and the resulting string (with the new-line character, if any, removed) is assigned to the variable.

The following are examples of simple assignments. You may omit quotes around the arguments if you are sure that they contain no special characters:

```
= a "$1"
= b '*****'
= c /usr/news/.mail
```

The procedure below illustrates the use of a default argument. If an argument is given, mail is read from it. Otherwise, mail is read from "/usr/news/.mail":

```
= a "$1" /usr/news/.mail
mail -f $a
```

The "=" command is often used to capture the output of a program. For example, *date* writes the current time and date to its standard output. The following line saves this value in \$d:

```
date | = d
```

This works just as well in longer pipelines. The following saves in \$a the number of logged-in users:

```
who | wc -l | = a
```

Another use is in the writing of *interactive* Shell procedures. The following example is part of a procedure to ask the user what kind of terminal is being used, so that tabs and delays can be set and other useful actions taken. The "</dev/tty" indicates a redirection of the standard input to the user terminal; it is *not* seen as an argument to "=", but rather causes the variable to be set to the next line typed by the user:

```
echo 'terminal?'
= a </dev/tty
```

Several variables are currently assigned special meanings:

\$n records the number of arguments passed to the Shell, not counting the name of the Shell procedure itself. Thus, "sh file arg1 arg2 arg3" sets \$n to 3. Its primary use is in checking for the required number of arguments:

```
if $n -lt 2 then
    echo 'two or more args required'; exit
endif
```

Shift never changes the value of \$n.

\$p permits alteration of the ordered list of directory pathnames used when searching for commands. It contains a sequence of directory names (separated by colons) that are to be used as search prefixes, ordered from left to right. The current directory is indicated by a null string.

By default, `$p` is initialized to a value producing the effect described in {3.7}: `"/bin:/usr/bin"`. A user could possess a personal directory of commands (say, `/a1/tf/jtb/bin`) and cause it to be searched *before* the other three directories by using:

```
= p /a1/tf/jtb/bin::/bin:/usr/bin
```

- `$r` gives the value of the return code of the command most recently executed by the Shell. It is a string of digits; most commands return "0" to indicate successful completion. For example, the "=" command returns "0" if two arguments are given and the first is not null, or if a line is actually read from the input. When the Shell terminates, it returns the current value of `$r` as its own return code.
- `$s` is initialized to the name of the user's *login directory*, i.e., the directory that becomes the current directory upon completion of a login (e.g., `"/a1/tf/jtb"`). Using this variable helps one to keep full pathnames out of Shell procedures. This is of great benefit when pathnames are changed, either to balance disk loads or to reflect administrative changes.
- `$t` is initialized to the user's terminal identification, a single letter or digit. The terminal can be manipulated using the file name `"/dev/tty$t"` or just `"/dev/tty"` alone. The latter is a generic name for the user's terminal.
- `$w` is initialized to the first component of `$s`, i.e., it is the name of the file system (such as `"/a1"`) in which the login directory is located. Like `$s`, it is used to avoid pathname dependencies, but is more useful than `$s` for projects involving many users.
- `$z` is initialized to `"/bin/sh"`. The command named by `$z` is the one that actually reads the Shell procedures invoked implicitly. The user can alter the choice of the Shell by overriding this value {4.4}. This facility is very useful when there are several different Shells in a system. This may occur because different groups of users want different Shells, or when a new Shell is being tested.

In addition to the above variables, the following read-only variable is provided:

- `$$` contains a 5-digit number that is the unique process number of the current Shell. Its most common use is in generating unique names for temporary files. Unlike many other systems, PWB/UNIX provides no separate mechanism for the automatic creation and deletion of temporary files: a file exists until it is explicitly removed. Temporary files are generally undesirable objects: the PWB/UNIX pipe mechanism is far superior for many applications. However, the need for uniquely-named temporary files does occur, especially for multi-user database applications. The following example of `$$` usage also illustrates the helpful practice of creating temporary files in a directory used only for that purpose:

```
ls >$s/tmp/$$
... commands (some of which use $s/tmp/$$)
: 'clean up at end'
rm $s/tmp/$$
```

4.4 Initialization of `$p` and `$z` by the `.path` File

The user may request automatic initialization of each Shell's `$p` (and `$z`) by creating a file named `".path"` in the login directory. The first (or only) line should be of the form shown for `$p` {4.3}. If present, the second line should be the full pathname of a Shell. Every instance of the Shell looks for that `".path"` file and initializes its own `$p` (and `$z`) from it, if `".path"` exists. Otherwise, `"/bin:/usr/bin"` and `"/bin/sh"` are the values used, respectively. Thus, the `".path"` information is available to all of the user's Shells, but changing `$p` or `$z` in one Shell does *not* affect these variables in other Shells. In addition, `".path"` is used in a consistent way by commands that must search for other commands, such as *nohup*, *nice*, and *time*.⁶ This facility is heavily used in large projects, because it simplifies the sharing of procedures, and can be quickly altered to adapt to changes in organizational requirements.

6. If you plan to write such a command, investigate the *pevec* subroutine, which combines the search and execution code

4.5 Control Structures

The Shell provides several commands that implement a variety of control structures. These commands are presented here in order of increasing complexity. See {6} for examples of these commands in the context of complete Shell procedures.

☛ Several of the control commands must not be "hidden" on command lines (e.g., behind semi-colons ";"):

else end endif endsw if switch while

Other control commands may be "hidden":

break breaksw continue exit goto next

4.5.1 Labels and Goto. The command ":" is recognized by the Shell, but is then treated as a null operation. One use of ":" is to define a label to act as a target for *goto*. Another use is to begin a comment line. However, it is a good idea to place comments in quotes {3.5} if they contain any characters that have a special meaning to the Shell, because the line *is* actually parsed, not just ignored. Using "goto label" causes the following actions:

- A *seek* is performed to move the read pointer to the beginning of the command file.
- The file is scanned from the beginning, searching for ": label", either alone on a line, or followed by a blank or tab.
- The read pointer is made to point at the line *after* the labeled line.

Thus, the only effect of *goto* is the adjustment of the Shell's file read pointer to cause the Shell to resume interpreting commands starting at the line following the labeled line. Invoking *goto* with an undefined label causes termination of the procedure {4.5.5}.

☛ Avoid the "goto"—future versions of the Shell are not expected to allow it.

4.5.2 If: Simple Conditional.

if conditional-expression command [args]

Whenever the conditional-expression is found to be *true*, *if* executes the command (via the *exec* system call), passing the arguments to it. Whenever the conditional-expression is *false*, *if* merely exits.

The following primaries can be used to construct the conditional-expression:

- r file *true* if the named file exists and is readable by the user.
- w file *true* if the named file exists and is writable by the user.
- s file *true* if the named file exists and has a size greater than zero.
- d file *true* if the named file is a directory.
- f file *true* if the named file is an ordinary file.
- s1 = s2 *true* if strings *s1* and *s2* are identical.
- s1 != s2 *true* if strings *s1* and *s2* are *not* identical.
- n1 -eq n2 *true* if the integers *n1* and *n2* are algebraically equal. Other algebraic comparisons are indicated by "-ne", "-gt", "-ge", "-lt", and "-le".
- { command } the command is executed; a return code of 0 (yes, zero!) is considered *true*, any other value is considered *false*. Most commands return 0 to indicate successful completion.

These primaries may be combined with the following operators:

- ! unary negation operator.
- a binary logical *and* operator.
- o binary logical *or* operator; it has lower precedence than "-a".
- (expr) parentheses for grouping. They must be escaped to remove their significance to the Shell. In the absence of parentheses, evaluation proceeds from left to right.

All of the operators, flags, and values are *separate* arguments to *if*, and must be separated by blanks. You must be careful to make sure that an argument actually appears and can be parsed correctly:

```
if "$1" = "" echo missing argument
if 0$1 = 0 echo missing argument
if 0"$1" = 0 echo missing argument
```

The first example guards against the possibility that \$1 is omitted, null, or has embedded blanks; the second guards against the possibility that \$1 has a value that causes parsing problems (such as “-r”), or that it is omitted or null; the third guards against all these problems. The following is dangerous:

```
if $1 = "" echo missing argument
```

because it would cause a syntax error in any of the above cases. Substitution of variables and arguments occurs effectively *before* parsing; thus, for example, if \$1 were null, then after substitution the line would read:

```
if = "" echo missing argument
```

In this case, \$1 without quotes yields no argument at all (on the other hand, "\$1" would have yielded an argument whose value is the null string). It is generally desirable to quote arguments (with double quotes—see 3.5), especially when they might possibly contain blanks or other characters that have a special meaning to the Shell. Examples of the use of *if* can be found in {6}.

4.5.3 If-then-else-endif: Structured Conditional. A more general (and much more readable) form of *if* can be used:

```
if conditional-expression then
    ... commands
else
    ... commands
endif
```

The *else* and the commands following it may be omitted. It is legal to nest *if* commands, but there must be an *endif* to match every *then*.

When *if* is called with a command, using the form of {4.5.2}, it acts as described there, deciding whether or not to execute the supplied command. When called with *then* instead of another command, *if* simply exits on a *true*, allowing the Shell to read and interpret the immediately following lines. On a *false*, *if* reads the file until it finds the next unmatched *else* or *endif*, thus skipping it and any other intervening lines. *Else* reads to the next unmatched *endif*. *Endif* is a null command.

These commands work together in a way that produces the appearance of a familiar control structure, although they do little but adjust the Shell's read pointer. Be warned that this implementation technique does *not* do a good job of diagnosing extra, missing, or hidden *if*, *else*, or *endif* commands {4.5}; if you suspect that there are such extra or missing commands, “opt -v” often helps {3.8,4.1}.

4.5.4 Switch-breaksw-endsw: Multi-way Branch. The *switch* command manipulates the input file in a way quite similar to *if*. It is modeled on the “switch” statement of the C language [8], and like it, provides an efficient multi-way branch:

```
switch value
: label1
    ... commands
: label2
    ... commands
:
: default
    ... commands
endsw
```

Switch reads the input until it finds:

- a statement label that matches *value*. The label may contain special characters as described in {3.4}; the method of matching is identical. A few of the many possible labels that could be used to match the value "thing.c" are:

```
thing.c  *.c  t*  *  ????????
```

- *default* used as a statement label (optional).
- the next unmatched *endsw* command.

Again, from the Shell's viewpoint, the only effect of *switch* is to adjust the read pointer so that the Shell effectively skips over part of the procedure, and then continues executing commands following the chosen label or *endsw*. For examples, see ".profile" and "fsplit" in {6}.

Value is obtained from an argument or from a variable; if the label *default* is present, it must be the last label in the list; it indicates a default action to be taken if *value* matches none of the preceding labels. The *switch* construct may be nested; labels enclosed by interior *switch-endsw* pairs are ignored during the execution of *switch*. *Breaksw* reads the input until the next unmatched *endsw* and is used to end the sequence of commands associated with a label. *Endsw* is a null command like *endif*.

4.5.5 End-of-file and Exit. When the Shell reaches the end-of-file, it terminates execution, returning to its parent the return code found in \$r. The *exit* command simply seeks to the end-of-file and returns, setting the return code to the value of its argument, if any. Thus, a procedure can be terminated "normally" by using *exit 0*.

4.5.6 While—break—continue—end: Looping. A *while-end* pair delimits a loop. *Break* can be used to terminate execution of such a loop. *Continue* requests the execution of the next iteration of the loop:

```
while conditional-expression
    ... commands
end
```

While evaluates the conditional-expression, which is similar to that of *if* {4.5.2}. If the conditional-expression is *true*, *while* does nothing, permitting the following lines to be read and interpreted. If the conditional-expression is *false*, the input file is searched for a matching *end*, and command interpretation resumes with the next line. *While-end* groupings may be nested to a depth of three.

While treats a single, non-null argument as *true* and a single null argument or lack of arguments as *false*. This is convenient for the simple case that handles one argument per iteration:

```
while "$1"
    Do something with $1.
    shift
end
```

Break terminates execution of the smallest enclosing *while-end* group, causing execution to resume after the nearest following unmatched *end*. Exit from *n* levels is obtained by writing *n* *break* commands on the same line:

```
break; break; ...
```

Continue causes execution to resume at the preceding *while*, i.e., the one that begins the smallest loop containing the *continue*.

4.5.7 Conditional Operators || and &&. These operators enforce left-to-right execution of commands. In the line "cmd1 || cmd2", *cmd1* is executed and its return code examined. Only if it failed (exit code non-zero) is *cmd2* executed. It is thus a more terse notation for:

```
cmd1
if $r -ne 0 then
    cmd2
endif
```

The “&&” operator yields the inverse test: in “cmd1 && cmd2”, the second command is executed only if the first succeeds (exit code zero). In the sequence below, each command is executed in order, until one fails:

```
cmd1 && cmd2 && cmd3 && ... && cmdn
```

See “fsplit” and “writemail” in {6} for examples.

4.5.8 Next: Transfer to Another File. The command “next name” causes the Shell to abandon the current input and begin reading file name. *Next* with no arguments causes the Shell to read from the terminal. By creating a file that initializes Shell variables, then typing “next file” at the terminal, anyone can have a simple shorthand for setting a number of Shell variables with little typing. See “nx” in {6}.

4.6 Onintr: Interrupt Handling

As noted in {2.2}, a program may choose to “catch” an interrupt from the terminal, ignore it completely, or be terminated by it. Shell procedures can use *onintr* to obtain the same effects:

```
onintr [ label ]
```

Onintr takes several forms: “onintr label” yields the effect of “goto label” on receipt of an interrupt; “onintr” alone causes normal action to be restored, so that the process terminates on the next interrupt; “onintr -” causes interrupts to be ignored completely, not only by the Shell, but also by any commands invoked by it.

The most frequent use of *onintr* is to make sure that temporary files are removed at the end of a procedure. The example at the end of {4.3} typically would be written as:

```
onintr clean
ls >$$/tmp/$$
... commands
: clean
rm $$/tmp/$$
```

When “onintr label” is used, interrupts are effective at the time when the label is reached; it is often desirable to insert another *onintr* following the label. Even so, there may be a short “window” when the user can accidentally kill the procedure by causing repeated interrupts in quick succession.

4.7 Special I/O Redirections

As noted in {3.2}, when the Shell is invoked it expects to inherit from its parent an open standard input (file descriptor 0), standard output (file descriptor 1), and diagnostic output (file descriptor 2). Each of these is initially connected to the terminal.

4.7.1 Standard Input. When the Shell is invoked to read a command file, it saves the old standard input (in an invisible place), then opens the command file as the new standard input. The fact that commands inherit the *new* standard input is convenient for commands that read in-line data (editor scripts, etc.) not read by the Shell. However, this mechanism prevents a Shell procedure from acting as a filter or from reading the *old* standard input in the way that most C programs do. The Shell solves this problem by permitting the notation “<--” to allow a command to take its input from the old standard input, which the Shell has previously saved.⁷

Note that “</dev/tty” and “<--” usually have equivalent effects in a procedure invoked directly from the terminal. The effects differ in a procedure invoked from within another procedure, unless the first procedure takes care to invoke the second with “<--”. In any case, “<--” is to be preferred because it can be used to read from a file or a pipe and is thus more general. See “fsplit” and “lower” in {6}.

7. The notation “<--” arises from the concept of “standard input once removed”, because many PWB/UNIX commands accept “<-” in place of a file name to indicate that the *current* standard input should be read. This choice makes it impossible to redirect input from a file named “<-”. Fortunately, file names almost never begin with “<-”, because many commands expect “<-” to signal a flag of some sort.

4.7.2 Standard Output. The use of “>/dev/tty” redirects output to the terminal, even if used in the middle of a pipeline. Shell procedures that act as filters sometimes need to do this. The redirection “>/dev/null” causes the standard output of a command to be thrown into a bottomless pit (presumably to feed the wumpus—see *wump(VI)*). This is used when you want to execute a command for its side-effects, but do not want to be bothered by its output.

4.7.3 Diagnostic Output. Most commands direct diagnostics to file descriptor 2 to make sure that they do not get lost down pipelines. Some situations require that this output go to some place other than the terminal. For example, a long-running procedure may be started, and then the terminal is hung up. In this case, it is helpful to save diagnostics in a file. A deficiency of the current Shell is the lack of syntax for redirecting the diagnostic output. The separate command *fd2* performs the required services:

```
fd2 [ + ] [ -file ] [ --file ] command arguments ...
```

The “+” flag causes diagnostic output to be merged into the standard output. The second option writes that output to *file*; the third appends it to *file*. If the file name is omitted in the second or third cases, “msg.out” is used. If no flag is given, “-msg.out” is assumed.

4.8 Quoting Revisited

The main problem with quoting conventions is the need to treat “\$” and “\” in ways flexible enough for convenient use with arguments and variables, but simple enough to be understandable, easy to implement, and unobtrusive in simple cases. In this respect, the current version of the Shell is far from elegant, but is reasonable in practice. The rules are:

- *Inside single quotes*, every character stands for itself without exception. A single quote is *not*, itself, allowed within single quotes.
- *Inside double quotes*, “\” and “\” stand for the characters “\$” and “””, respectively, but with all special meaning removed. All other characters, other than a pair of characters the first of which is an unescaped “\$”, behave exactly as they do within single quotes, including a “\” *not* followed by a “\$” or a “””.
- *Inside double quotes and outside either kind of quotes*, any two-character sequence whose first character is an unescaped “\$” is replaced by the value of the corresponding Shell argument or variable; any variable that has no value (such as “\$:”) is replaced by a null string.
- *Outside either kind of quotes*, any two-character sequence whose first character is a “\” is replaced by the second character of that sequence, but with any special meaning removed.

4.9 Creation and Organization of Shell Procedures

A Shell procedure can be created in two simple steps. The first is that of building an ordinary text file. The second is that of changing the *mode* of the file to make it *executable*, thus permitting it to be invoked by “name args”, rather than “sh name args”. The second step may be omitted for a procedure to be used once or twice and then discarded, but is recommended for longer-lived ones.

Here is the entire input needed to set up a simple procedure (the executable part of “draft” in {6}):

```
ed
a
nroff -rC3 -T450-12 -mm $*
.
w draft
q
chmod 755 draft
```

It may then be invoked as “draft file1 file2”. If the Shell procedure “draft” were thus created in a directory whose name appears in the user’s “.path” file, the user could change working directories and still invoke the “draft” command.

Shell procedures may be created dynamically. A procedure may generate a file of commands, invoke another instance of the Shell to execute that file, then remove it. An alternate approach is that of using *next* to make the current Shell execute the new file, allowing use of existing Shell variables and avoiding the spawning of an additional process for another Shell. In some cases, the need for a temporary file may be eliminated by using the Shell in a pipeline.

Many users prefer to write Shell procedures instead of C programs. First, it is easy to create and maintain a Shell procedure because it is only an ordinary file of text. Second, it has no corresponding object program that must be generated and maintained. Third, it is easy to create a procedure “on the fly”, use it a few times, then remove it. Finally, because Shell procedures are usually short in length, written in a high-level programming language, and kept only in their source-language form, they are generally easy to find, understand, and modify.

By convention, directories of commands and/or Shell procedures are usually named “bin”. Most groups of users sharing common interests have one or more “bin” directories set up to hold common procedures. Some users have “.path” files that list several such directories. Although you can have a number of such directories, it is unwise to go overboard—it may become difficult to keep track of your environment, and efficiency may suffer {7.3}.

5. MISCELLANEOUS SUPPORTING COMMANDS

Shell procedures can make use of almost any command. The commands described in this section are either used especially frequently in Shell procedures, or are explicitly designed for such use.

5.1 Echo: Simple Output

The *echo* command, invoked as “echo [args]”, copies its arguments to the standard output, each followed by a single space, except the last argument, which is followed by a new-line; often, it is used to prompt the user for input, to issue diagnostics in Shell procedures, or to add a few lines to an output stream in the middle of a pipeline. Another use is to verify the argument list generation process (as in {3.4}) before issuing a command that does something drastic. The command “ls” is often replaced by “echo *” because the latter is faster and prints fewer lines of output.

Echo recognizes several escape sequences. A “\n” yields a new-line character. *Echo* normally appends a new-line character to its last argument; a “\c” is used to *suppress* that new-line character. The following prompts the user for input and allows input to be typed on the same line as the prompt:

```
echo 'enter name: \c'  
= a </dev/tty
```

Echo also recognizes an octal escape sequence for any character, whether printable or not.

5.2 Pump: Shell Data Transfer

Pump is a filter that copies its standard input to its standard output with possible substitution of Shell arguments and variables:

```
pump [ -[ subchar ] ] [ + ] [ eofstr ]
```

Pump reads input until an end-of-file, or until it finds *eofstr* alone on a line. The default *eofstr* is “!”. Normally, Shell arguments and variables are substituted in the data stream. The flag “-” suppresses all substitution, while the form “-subchar” causes *subchar* to be used as the indicator character for substitution of Shell variables and arguments, instead of “\$”. Escaping is handled as in strings enclosed by double quotes—the indicator character may be hidden by preceding it with “\”. The “+” flag causes all leading tab characters in the input to *pump* to be eliminated; this permits that input to be indented for readability. A common use of *pump* is to get Shell variables into editor scripts—see “edfind” in {6}, for example. Because editor scripts may use “\$” for other purposes, readability may be improved by using a *subchar* such as “%”:

```
:      'in file $1, change every instance of $2 to $3'
:      'then delete all lines consisting only of $4'
if -r "$1" then
    pump -% + | ed $1
    g/%2/s//%3/g
    g/^%4$/d
    w
    !
else
    echo "$1: cannot open"
endif
```

Pump is often used to copy a few lines to another file:

```
pump >>logfile
here is $1
and here is $2 on a separate line
!
```

5.3 Expr: Expression Evaluation

Expr supports arithmetic and logical operations on integers, and PL/I-like “substr”, “length”, and “index” operators for string manipulation. It evaluates a single expression and writes the result to the standard output, typically piped into “=” to be assigned to a variable. Typical examples are:

```
:      'increment $a'
expr $a + 1 | = a

:      'put 3rd through last characters of $1 into $b'
:      'expr substr abcde 3 1000 returns cde (1000 is just a big number)'
expr substr "$1" 3 1000 | = b

:      'obtain length of $1'
expr length "$1" | = c
```

The most common uses of *expr* are in counting for loops and in using “substr” to pick apart strings.

5.4 Logname, Logdir, Logtty: Login Data

When a user logs in, he or she supplies a *login name* and a password. The *login* program searches the password file for that *login name* and obtains the name of the program to be executed by the user (normally the Shell), the directory to be made the current directory, and also a *userid*, a value ranging from 0 to 255. Most UNIX protection and identification mechanisms utilize the last item. Limiting the number of distinct users to 256 is no problem for most UNIX systems, but the original PWB/UNIX installation currently supports more than 1,000 users. However, it is not necessary to provide a distinct userid for every user. Project-oriented groups of users often choose to share one or two userids, in order to ease the problems caused by personnel absences, and also to ease the manipulation of shared files.⁸ Although the members of such groups do not generally worry about being protected from each other, they need to be identified as distinct individuals by some programs, i.e., those that tag inter-user messages with user names or log the name of the user making a change to a source program. PWB/UNIX records the login name instead of discarding it after login. The *logname* command writes this name to the standard output, allowing it to be captured in a Shell variable. It can then be used to permit only selected users to execute a procedure, or can be included in logging information:

```
logname | = u
(echo "$u updated files on \c"; date) >>projectlog
```

The *logdir* and *logtty* commands are used in the same way as *logname*; they produce the same values as the initial values of \$s and \$t, respectively {4.3}.

8. Although some groups started by using one userid per person, it was discovered that these users often shared a single password. Thus, the possession of separate userids was considered more of a hindrance than a help.

6. EXAMPLES OF SHELL PROCEDURES

☛ *Some examples in this section may be quite difficult for beginners. For ease of reference, the examples are arranged alphabetically by name.*

.profile:

```
:      '.profile (automatically invoked on login) asks for terminal type,'
:      'reads a line from terminal, loops until a known type'
:      '(or empty line) is entered, sets terminal options appropriately,'
:      'asks for new directory name and changes to it, if one is given,'
:      'and then, if file nx exists, transfers to it'
while 1
  echo 'terminal:\c'
  = a </dev/tty
  switch "$a"
  :      'DASI450'
  :      450
  :      stty cr2; tabs +t450; break
  :      'GSI/DASI300'
  :      gsi
  :      300
  :      stty cr2; tabs; break
  :      'HP264X'
  :      hp
  :      stty cr0 nl0; tabs +thp; break
  :      'TI 700'
  :      ti
  :      stty -tabs nll cr1; break
  :      default
  :      if 0"$a" = 0 break
  :      echo "$a? try 450,gsi,hp,ti"
  endsw
end
echo "cd \c"
= b </dev/tty
if "$b" != "" then
  cd $b
endif
if -r nx then
  next nx
endif
```

Note: *Break* is used instead of *breaksw* in the above example to terminate the *while* loop, not just the *switch* construct.

copypairs:

```
:      'copypairs file1 file2 ...'
:      'copy file1 to file2, file3 to file4, ...'
while "$2"
  cp $1 $2
  shift; shift
end
if 0"$1" != 0 echo 'odd number of arguments'
```

Note: Remember that "shift; shift" is *not* the same as "shift 2". See next example for use of "shift 2".

copyto:

```
:      'copyto dir file ...
:      'copy argument files to dir, making sure that at least'
:      'two arguments exist, that dir is a directory, and that'
:      'each additional argument is a readable file'
if $n -lt 2 then
    echo 'usage: copyto directory file ...'; exit
endif
if ! -d $1 then
    echo "$1 is not a directory"; exit
endif
while "$2"
    if ! -r $2 then
        echo "$2 not readable"
    else
        cp $2 $1
    endif
    shift 2
end
```

distinct1:

```
:      'distinct1'
:      'reads standard input, reports list of identifiers that'
:      'differ only in case, giving lower case form of each'
tr -cs '[A-Z][a-z][0-9]' '\012*' <-- | sort -u | tr '[A-Z]' '[a-z]' | sort | uniq -d
```

Note: This procedure is an example of the kind of process that is created by the "left-to-right" construction of a long pipeline. It may not be immediately obvious how this works. The *tr* translates all characters except letters and digits into new-line characters, and then "squeezes out" repeated new-line characters. This leaves each identifier (in this case, any contiguous sequence of letters and digits) on a separate line. *Sort* sorts the lines and emits only one line from any sequence of one or more repeated lines. The next *tr* converts everything to lower case, so that identifiers differing only in case become identical. The output is sorted again to bring such duplicates together. The *uniq -d* prints once only those lines that occur more than once, yielding the desired list.

The process of building such a pipeline uses the fact that pipes and files can usually be interchanged; the two lines below are equivalent, assuming that sufficient disk space is available:

```
cmd1 | cmd2 | cmd3
cmd1 >tmp1; <tmp1 cmd2 >tmp2; <tmp2 cmd3; rm tmp[1-3]
```

Starting with a file of test data and working from left to right, each command is run taking its input from the previous file and putting its output in the next file. The final output file is then examined to make sure that it contains the expected result. The goal is to create a series of transformations that will convert the input to the desired output. As an exercise, try to mimic "distinct1" with such a step-by-step process, using a file of test data containing:

```
ABC:DEF/DEF
ABC1 ABC
Abc abc
```

Although pipelines can give a concise notation for complex processes, exercise some restraint lest you succumb to the "one-line syndrome" sometimes found among users of especially concise languages. This syndrome often yields incomprehensible code.

distinct2:

```
:      'distinct2'
:      'reads standard input, reports sorted list of identifiers that differ'
:      'in case only, listing all such distinct identifiers'
onintr cleanup
tr -cs '[A-Z][a-z][0-9]' '[\012*]' <-- | sort -u | tee t1$$ | tr '[A-Z]' '[a-z]' >t2$$
pr -s -t -l1 -m t1$$ t2$$ | sort +1 >t3$$
:      'third argument to pr in above line is "minus ell one"'
sort t3$$ >t4$$
uniq -u -l t3$$ | sort | comm -23 t4$$ - | sort +1
: cleanup
rm t?$$
```

Note: This procedure is similar to the previous one, but provides more explicit information. As an exercise, work through this procedure in the way described above. The commands used here (plus *grep* and *sed*) form the basis for many "data stream" operations.

draft:

```
:      'draft file ...'
:      'prints the draft (-rC3) of a document on a DASI450 terminal in 12-pitch'
:      'using PWB/MM'
nroff -rC3 -T450-12 -mm $*
```

Note: Users often write this kind of procedure for convenience in dealing with commands that require the use of many distinct flags that cannot be given default values that are reasonable for all (or even most) users.

edfind:

```
:      'edfind file arg'
:      'find the last occurrence in file of a line that matches arg,'
:      'then print 3 lines (the one before, the line itself, and the one after)'
pump | ed - $1
?S2?;-,+p
!
```

Note: This illustrates the typical practice of using *pump* to substitute Shell variables into *ed* scripts.

edlast:

```
:      'edlast file'
:      'prints the last line of file, then deletes that line'
ed - $1
$p
$d
w
q
echo done
```

Note: This procedure illustrates the effects of a command that reads input from a file shared with the Shell.

fsplit:

```
: 'fsplit file1 file2'
: 'read standard input and split it into three parts:'
: 'append any line containing at least one letter to file1, any line'
: 'containing digits but no letters to file2, and throw the rest away'
= i 0; = j 0
while 1
  = a <-- || break
  expr $i + 1 | = i
  switch "$a"
  : *[A-Za-z]*
    echo "$a" >>$1; breaksw
  : *[0-9]*
    echo "$a" >>$2; breaksw
  : default
    expr $j + 1 | = j
  endsw
end
echo "$i lines read, $j thrown away"
```

Note: Each iteration of the loop reads a line from the input and analyzes it. The *break* terminates the loop only when "=" encounters an end-of-file.

☛ *Don't use the Shell to read a line at a time unless you must—it can be grotesquely slow {7.2.1}.*

loop:

```
: 'loop arg ...'
: 'one or more command lines'
: 'endloop'
: 'execute the group of command lines once for each argument,'
: 'substituting each argument as $1 in the command lines'
onintr cleanup
echo 'while "$1" >tmp$$
pump - + endloop <-- >>tmp$$
echo 'shift \n end' >>tmp$$
next tmp$$; rm tmp$$
: cleanup
rm tmp$$
```

Note: Such a procedure is typically used from a terminal to repeat some commands for a list of arguments. It creates a temporary file that sandwiches user input between a *while* and *shift-end*. It then transfers to that file. For example, all files in the current directory could be copied to "place" by:

```
loop *
cp $1 place
echo $1 copied
endloop
```

lower:

```
: 'lower'
: 'reads standard input, converts it to lower case, writes to standard output'
: 'can thus be used in a pipeline if desired'
tr '[A-Z]' '[a-z]' <--
```

Note: This is the most common type of use for "<--".

mkfiles:

```
:      'mkfiles prefix [number]'
:      'makes number (default = 5) files, named prefix1, prefix2, ...'
= a "$2" 5
= i 1
while $i -le $a
    cp /dev/null $1$i
    expr $i + 1 | = i
end
```

null:

```
:      'null file ...'
:      'create each of the named files as an empty file'
while "$1"
    cp /dev/null $1
    shift
end
```

nx:

```
:      'next nx'
:      'asks for module name, initializes variables to useful values,'
:      'prints variables. Note that variables are set within the invoking Shell,'
:      'so nx can be invoked only from terminal or from .profile'
= a /sys/source/sl
= b /usr/man/man1
echo "m: \c"
= m </dev/tty
= g "get -e s.$m; ed $m"
= d "delta s.$m"
pump
a: $a   b: $b
d: $d   g: $g   m: $m
!
next
```

phone:

```
:      'phone initials'
:      'prints the phone number(s) of person with given initials'
echo 'inits   ext   home'
grep "^$1"
abc    1234   999-2345
def    2234   583-2245
ghi    3342   988-1010
xyz    4567   555-1234
```

writemail:

```
:      'writemail message user'
:      'if that user is logged in, write message on terminal;'
:      'otherwise, mail it to that user'
echo "$1" | ( write "$2" || mail "$2" )
```

Note: Replacing "echo" above by "pump . <--" writes or mails the standard input, in the same way as the *mail* command.

7. EFFECTIVE AND EFFICIENT SHELL PROGRAMMING

7.1 Overall Approach

This section outlines strategies for writing "efficient" Shell procedures, i.e., ones that do not waste resources unreasonably in accomplishing their purposes. In the author's opinion, the primary reason for choosing the Shell procedure as the implementation method is to achieve a desired result at a minimum *human* cost. Emphasis should *always* be placed on simplicity, clarity, and readability, but efficiency can also be gained through awareness of a few design strategies. In many cases, an effective redesign of an existing procedure improves its efficiency by reducing its size, and often increases its comprehensibility. In any case, one should not worry about optimizing procedures unless they are intolerably slow or are known to consume a lot of resources.

The same kind of iteration cycle should be applied to Shell procedures as to other programs: write code, measure it, and optimize only the *few* important parts. The user should become familiar with the *time* command, which can be used to measure both entire procedures and parts thereof. Its use is strongly recommended; human intuition is notoriously unreliable when used to estimate timings of programs, even when the style of programming is a familiar one. Each timing test should be run several times, because the results are easily disturbed by, for instance, variations in system load.

7.2 Approximate Measures of Resource Consumption

7.2.1 Number of Processes Generated. When large numbers of short commands are executed, the actual execution time of the commands may well be dominated by the overhead of spawning processes. The CPU overhead per process lies in the range of 0.07 to 0.1 seconds, depending on the specific hardware configuration. The procedures that incur significant amounts of such overhead are those that perform much looping, and those that generate command sequences to be interpreted by another Shell.

If you are worried about efficiency, it is important to know which commands are currently built into the Shell, and which are not. Here is the alphabetical list of those that are built-in:

:	chdir	endsw	newgrp	shift
=	continue	exit	next	switch
break	else	goto	onintr	test
breaksw	end	if	opt	wait
cd	endif	login	pump	while

Pump actually executes as a child process, i.e., the Shell does a *fork*, but no *exec*; "(" executes in the same way. Any command *not* in the above list requires both *fork* and *exec*.

The user should always have at least a vague idea of the number of processes generated. In the bulk of observed procedures, the number of processes spawned (not necessarily simultaneously) can be described by:

$$\text{processes} = k * n + c$$

where *k* and *c* are constants, and *n* is the number of procedure arguments, the number of lines in some input file, the number of entries in some directory, or some other obvious quantity. Efficiency improvements are most commonly gained by reducing the value of *k*, sometimes to zero. Any procedures whose complexity measures include *n*² terms or higher powers of *n* are likely to be intolerably expensive.

As an example, here is an analysis of procedure "fsplit" of {6}. For each iteration of the loop, there is one *expr* plus either an *echo* or another *expr*. One additional *echo* is executed at the end. If *n* is the number of lines of input, the number of processes is 2 * *n* + 1. On the other hand, the number of processes in the following (equivalent) procedure is 12, regardless of the number of lines of input:

fsplit2:

```
onintr cleanup
= b '[ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz]'
cat <-- >tmp$$
grep "$b" tmp$$ >tmp$$1
grep -v "$b" tmp$$ | grep "[0123456789]" >tmp$$2
cat tmp$$1 >>$1 ; cat tmp$$2 >>$2
wc -l <tmp$$ | = i
wc -l <tmp$$1 | = j
wc -l <tmp$$2 | = k
expr $i - $j - $k | = a
echo "$i read, $a thrown away"
:      cleanup
rm tmp$$*
```

This version is often ten times faster than “fsplit”, and it is even better for larger input files.

Some types of procedures should *not* be written using the Shell. For example, if one or more processes are generated for each character in some file, it is a good indication that the procedure should be rewritten in C.

■ *Shell procedures should not be used to scan or build files a character at a time.*

7.2.2 Number of Bytes of Data Accessed. It is worthwhile considering any action that reduces the number of bytes read or written. This may be important for those procedures whose time is spent passing data around among a few processes, rather than creating large numbers of short processes. Some filters shrink their output, others usually increase it. It always pays to put the “shrinkers” first when the order is irrelevant. Which of the following is likely to be faster?

```
sort file | grep pattern
grep pattern file | sort
```

7.2.3 Directory Searches. Directory searching can consume a great deal of time, especially in those applications that utilize deep directory structures and long pathnames. Judicious use of *cd* can help shorten long pathnames and thus reduce the number of directory searches needed. As an exercise, try the following commands (on a fairly quiet system).⁹

```
time sh -c 'ls -l /usr/bin/* >/dev/null'
time sh -c 'cd /usr/bin; ls -l * >/dev/null'
```

7.3 Efficient Organization

7.3.1 Directory Search Order and the .path File. The “.path” file is a popular and convenient mechanism for organizing and sharing procedures. However, it must be used in a sensible fashion, or the result may be a great increase in system overhead that occurs in a subtle, but avoidable way.

The process of finding a command involves reading every directory included in every pathname that precedes the needed pathname in the current \$p variable. As an example, consider the effect of invoking *nroff* (/usr/bin/nroff) when \$p is “:/bin:/usr/bin”. The sequence of directories read is: “.”, “/”, “/bin”, “/”, “/usr”, and “/usr/bin”, i.e., a total of six directories. A long “.path” can increase this number significantly.

The vast majority of command executions are of commands found in “/bin” and, to a lesser extent, in “/usr/bin”. Careless “.path” setup may lead to a great deal of unnecessary searching. The following four examples are ordered from worst to best (at least with regard to efficiency):

⁹ You may have to do some reading in the *PWB/UNIX User's Manual* [3] to understand exactly what is going on in these examples.

```
:/a1/tf/jtb/bin:/a1/tf/bin:/bin:/usr/bin  
:/bin:/a1/tf/jtb/bin:/a1/tf/bin:/usr/bin  
:/bin:/usr/bin:/a1/tf/jtb/bin:/a1/tf/bin  
/bin::/usr/bin:/a1/tf/jtb/bin:/a1/tf/bin
```

The first one above should be avoided. The others are acceptable—choice among them is dictated by the rate of change in the set of commands kept in “/bin” and “/usr/bin”.

A procedure that is expensive because it invokes many short-lived commands may often be speeded up by changing \$p to resemble the last of the above four examples.

7.3.2 Good Ways to Set up Directories. It is wise to avoid directories that are larger than necessary. You should be aware of several “magic sizes”. A directory that contains entries for up to 30 files (plus the required “.” and “..”) fits in a single disk block and can be searched very efficiently. One that has up to 254 entries is still a “small” file; anything larger is usually a disaster when used as a working directory. It is especially important to keep login directories small, preferably one block at most.


ACKNOWLEDGMENTS


The Shell was originally written by K. Thompson; its basic structure has remained unchanged since then, although many features (and some warts!) have been added. The PWB/UNIX extensions were added by R. C. Haight, A. L. Glasser, and the author. Some constructs have been derived from similar ones in the recent Shell written by S. R. Bourne. A number of colleagues provided helpful comments during the writing of this tutorial; T. A. Dolotta, in addition, provided a great deal of editorial assistance. Finally, many thanks must go to the PWB/UNIX user community, and especially M. H. Bianchi and J. T. Burgess, who provided many suggestions and examples.

REFERENCES

- [1] Bianchi, M. H., and Wood, J. L. A User's Viewpoint on the Programmer's Workbench. *Proc. Second Int. Conf. on Software Engineering*, pp. 193-99, Oct. 13-15, 1976.
- [2] Dolotta, T. A., and Mashey, J. R. An Introduction to the Programmer's Workbench. *Proc. Second Int. Conf. on Software Engineering*, pp. 164-68, Oct. 13-15, 1976.
- [3] Dolotta, T. A., Haight, R. C., and Piskorik, E. M., eds. *PWB/UNIX User's Manual—Edition 1.0*. Bell Laboratories, May 1977.
- [4] Kernighan, B. W., and Plauger, P. J. Software Tools. *Proc. First National Conference on Software Engineering*, pp. 8-13, Sept. 11-12, 1975.
- [5] Kernighan, B. W., and Plauger, P. J. *Software Tools*. Reading, MA: Addison-Wesley, 1976.
- [6] Mashey, J. R. Using a Command Language as a High-Level Programming Language. *Proc. Second Int. Conf. on Software Engineering*, pp. 169-76, Oct. 13-15, 1976.
- [7] Mashey, J. R. PWB/UNIX Documentation Roadmap. Bell Laboratories, 1977.
- [8] Ritchie, D. M. C Reference Manual. Bell Laboratories, 1977.
- [9] Ritchie, D. M., and Thompson, K. The UNIX Time-Sharing System. *Comm. ACM* 17(7):365-75, July 1974.
- [10] Thompson, K. The UNIX Command Language. In *Structured Programming—Infotech State of the Art Report*, pp. 375-84. Infotech International Limited, Nicholson House, Maidenhead, Berkshire, England, 1976.
- [11] Thompson, K., and Ritchie, D. M. *UNIX Programmer's Manual—Sixth Edition*. Bell Laboratories, May 1975.







UNIX For Beginners

Brian W. Kernighan

Bell Laboratories,
Murray Hill, New Jersey 07974

ABSTRACT

This paper is meant to help new users get started on UNIX. It covers:

- basics needed for day-to-day use of the system — typing commands, correcting typing mistakes, logging in and out, mail, inter-console communication, the file system, printing files, redirecting I/O, pipes, and the shell.
- document preparation — a brief tutorial on the ROFF formatter for beginners, hints on preparing documents, and capsule descriptions of some supporting software.
- UNIX programming — using the editor, programming the shell, programming in C, other languages.

There is also an annotated UNIX bibliography.



UNIX for Beginners

Brian W. Kernighan

Bell Laboratories, Murray Hill, N. J.

In many ways, UNIX is the state of the art in computer operating systems. From the user's point of view, it is easy to learn and use, and presents few of the usual impediments to getting the job done.

It is hard, however, for the beginner to know where to start, and how to make the best use of the facilities available. The purpose of this introduction is to point out high spots for new users, so they can get used to the main ideas of UNIX and start making good use of it quickly.

This paper is not an attempt to re-write the *UNIX Programmer's Manual*; often the discussion of something is simply "read section x in the manual." (This implies that you will need a copy of the *UNIX Programmer's Manual*.) Rather it suggests in what order to read the manual, and it collects together things that are stated only indirectly in the manual.

There are five sections:

1. **Getting Started:** How to log in to a UNIX, how to type, what to do about mistakes in typing, how to log out. Some of this is dependent on which UNIX you log into (phone numbers, for example) and what terminal you use, so this section must necessarily be supplemented by local information.
2. **Day-to-day Use:** Things you need every day to use UNIX effectively: generally useful commands; the file system.
3. **Document Preparation:** Preparing manuscripts is one of the most common uses for UNIX. This section contains advice, but not extensive instructions on any of the formatting programs.
4. **Writing Programs:** UNIX is an excellent vehicle for developing programs. This section talks about some of the tools, but again is not a tutorial in any of the programming languages that UNIX provides.
5. **A UNIX Reading List.** An annotated bibliography of documents worth reading by new users.

I. GETTING STARTED

Logging In

Most of the details about logging in are in the manual section called "How to Get Started" (pages iv-v in the 5th Edition). Here are a couple of extra warnings.

You must have a UNIX login name, which you can get from whoever administers your system. You also need to know the phone number. UNIX is capable of dealing with a variety of terminals: Terminus 300's; Execuport, TI and similar portables; video terminals; GSI's; and even the venerable Teletype in its various forms. But note: UNIX will not handle IBM 2741 terminals and their derivatives (e.g., some Anderson-Jacobsons, Novar). Furthermore, UNIX is strongly oriented towards devices with *lower case*. If your terminal produces only upper case (e.g., model 33 Teletype), life will be so difficult that you should look for another terminal.

Be sure to set the switches appropriately on your device: speed (if it's variable) to 30 characters per second, lower case, full duplex, even parity, and any others that local wisdom advises. Establish a connection using whatever magic is needed for your terminal. UNIX should type "login:" at you. If it types *garbage*, you may be at the wrong speed; push the 'break' or 'interrupt' key once. If that fails to produce a login message, consult a guru.

When you get a "login:" message, type your login name *in lower case*. Follow it by a RETURN if the terminal has one. If a password is required, you will be asked for it, and (if possible) printing will be turned off while you type it, again followed by a RETURN (On M37 Teletypes always use NEWLINE or LINEFEED in place of RETURN).

The culmination of your login efforts is a percent sign "%". The percent sign means that UNIX is ready to accept commands from the terminal. (You may also get a message of the day just before the percent sign or a notification that you have mail.)

Typing Commands

Once you've seen the percent sign, you can type commands, which are requests that UNIX do something. Try typing

```
date
```

followed by RETURN. You should get back something like

```
Sun Sep 22 10:52:29 EDT 1974
```

Don't forget the RETURN after the command, or nothing will happen. If you think you're being ignored, type a RETURN; something should happen. We won't show the carriage returns, but they have to be there.

Another command you might try is `who`, which tells you everyone who is currently logged in:

```
who
```

gives something like

```
pjp    ttyf    Sep 22 09:40
bwk    ttyg    Sep 22 09:48
mel    ttyh    Sep 22 09:58
```

The time is when the user logged in.

If you make a mistake typing the command name, UNIX will tell you. For example, if you type

```
whom
```

you will be told

```
whom: not found
```

Strange Terminal Behavior

Sometimes you can get into a state where your terminal acts strangely. For example, each letter may be typed twice, or the RETURN may not cause a line feed. You can often fix this by logging out and logging back in. Or you can read the description of the command `stty` in section I of the manual. This will also tell you how to get intelligent treatment of tab characters (which are much used in UNIX) if your terminal doesn't have tabs. If it does have computer-settable tabs, the command `tabs` will set the stops correctly for you.

Mistakes In Typing

If you make a typing mistake, and see it before the carriage return has been typed, there are two ways to recover. The sharp-character `"#"` erases the last character typed; in fact successive uses of `"#"` erase characters back to the beginning of the line (but not beyond). So if

you type badly, you can correct as you go:

```
dd#atte##e
```

is the same as `"date"`.

The at-sign `"@"` erases all of the characters typed so far on the current input line, so if the line is irretrievably fouled up, type an `"@"` and start over (on the same line!).

What if you must enter a sharp or at-sign as part of the text? If you precede either `"#"` or `"@"` by a backslash `"\"`, it loses its erase meaning. This implies that to erase a backslash, you have to type two sharps or two at-signs. The backslash is used extensively in UNIX to indicate that the following character is in some way special.

Readahead

UNIX has full readahead, which means that you can type as fast as you want, whenever you want, even when some command is typing at you. If you type during output, your input characters will appear intermixed with the output characters, but they will be stored away by UNIX and interpreted in the correct order. So you can type two commands one after another without waiting for the first to finish or even begin.

Stopping a Program

You can stop most programs by typing the character `"DEL"` (perhaps called `"delete"` or `"rubout"` on your terminal). There are exceptions, like the text editor, where `DEL` stops whatever the program is doing but leaves you in that program. You can also just hang up the phone. The `"interrupt"` or `"break"` key found on most terminals has no effect.

Logging Out

The easiest way to log out is to hang up the phone. You can also type

```
login name-of-new-user
```

and let someone else use the terminal you were on. It is not sufficient just to turn off the terminal. UNIX has no time-out mechanism, so you'll be there forever unless you hang up.

Mail

When you log in, you may sometimes get the message

```
You have mail.
```

UNIX provides a postal system so you can send and receive letters from other users of the system. To read your mail, issue the command

mail

Your mail will be printed, and then you will be asked

Save?

If you do want to save the mail, type *y*, for "yes"; any other response means "no".

How do you send mail to someone else? Suppose it is to go to "joe" (assuming "joe" is someone's login name). The easiest way is this:

mail joe
*now type in the text of the letter
on as many lines as you like ...
after the last line of the letter
type the character "control-d",
that is, hold down "control" and type
a letter "d".*

And that's it. The "control-d" sequence, usually called "EOT", is used throughout UNIX to mark the end of input from a terminal, so you might as well get used to it.

There are other ways to send mail — you can send a previously prepared letter, and you can mail to a number of people all at once. For more details see **mail (1)**.

The notation **mail (1)** means the command **mail** in section (1) of the *UNIX Programmer's Manual*.

Writing to other users

At some point in your UNIX career, out of the blue will come a message like

Message from joe...

accompanied by a startling beep. It means that Joe wants to talk to you, but unless you take explicit action you won't be able to talk back. To respond, type the command

write joe

This establishes a two-way communication path. Now whatever Joe types on his terminal will appear on yours and vice versa. The path is slow, rather like talking to the moon. (If you are in the middle of something, you have to get to a state where you can type a command. Normally, whatever program you are running has to terminate or be terminated. If you're editing, you can escape temporarily from the editor — read the manual.)

A protocol is needed to keep what you type from getting garbled up with what Joe types. Typically it's like this:

Joe types "write smith" and waits.

Smith types "write joe" and waits.

Joe now types his message (as many lines as he likes). When he's ready for a reply, he signals it by typing (o), which stands for "over".

Now Smith types a reply, also terminated by (o).

This cycle repeats until someone gets tired; he then signals his intent to quit with (o+o), for "over and out".

To terminate the conversation, each side must type a "control-d" character alone on a line. ("Delete" also works.) When the other person types his "control-d", you will get the message "EOT" on your terminal.

If you write to someone who isn't logged in, or who doesn't want to be disturbed, you'll be told. If the target is logged in but doesn't answer after a decent interval, simply type "control-d".

On-line Manual

The UNIX Programmer's Manual is typically kept on-line. If you get stuck on something, and can't find an expert to assist you, you can print on your terminal some manual section that might help. It's also useful for getting the most up-to-date information on a command. To print a manual section, type "man section-name". Thus to read up on the **who** command, type

man who

If the section in question isn't in part I of the manual, you have to give the section number as well, as in

man 6 chess

Of course you're out of luck if you can't remember the section name.

II. DAY-TO-DAY USE

Creating Files — The Editor

If we have to type a paper or a letter or a program, how do we get the information stored in the machine? Most of these tasks are done with the UNIX "text editor" **ed**. Since **ed** is thoroughly documented in **ed (1)** and explained in *A Tutorial Introduction to the UNIX Text Editor*, we won't spend any time here describing how to use it. All we want it for right now is to make some *files*. (A file is just a collection of information stored in the machine, a simplistic but adequate definition.)

To create a file with some text in it, do the following:

```
ed      (invokes the text editor)
a      (command to "ed", to add text)
now type in
whatever text you want ...
        (signals the end of adding text)
```

At this point we could do various editing operations on the text we typed in, such as correcting spelling mistakes, rearranging paragraphs and the like. Finally, we write the information we have typed into a file with the editor command "w":

```
w junk
```

ed will respond with the number of characters it wrote into the file called "junk".

Suppose we now add a few more lines with "a", terminate them with ".", and write the whole thing out as "temp", using

```
w temp
```

We should now have two files, a smaller one called "junk" and a bigger one (bigger by the extra lines) called "temp". Type a "q" to quit the editor.

What files are out there?

The ls (for "list") command lists the names (not contents) of any of the files that UNIX knows about. If we type

```
ls
```

the response will be

```
junk
temp
```

which are indeed our two files. They are sorted into alphabetical order automatically, but other variations are possible. For example, if we add the optional argument "-t",

```
ls -t
```

lists them in the order in which they were last changed, most recent first. The "-l" option gives a "long" listing:

```
ls -l
```

will produce something like

```
-rw-rw-rw- 1 bwk  41 Sep 22 12:56 junk
-rw-rw-rw- 1 bwk  78 Sep 22 12:57 temp
```

The date and time are of the last change to the file. The 41 and 78 are the number of characters (you got the same thing from ed). "bwk" is the owner of the file — the person who created it.

The "-rw-rw-rw-" tells who has permission to read and write the file, in this case everyone.

Options can be combined: "ls -lt" would give the same thing, but sorted into time order. You can also name the files you're interested in, and ls will list the information about them only. More details can be found in ls(1).

It is generally true of UNIX programs that "flag" arguments like "-t" precede filename arguments.

Printing Files

Now that you've got a file of text, how do you print it so people can look at it? There are a host of programs that do that, probably more than are needed.

One simple thing is to use the editor, since printing is often done just before making changes anyway. You can say

```
ed junk
l,$p
```

ed will reply with the count of the characters in "junk" and then print all the lines in the file. After you learn how to use the editor, you can be selective about the parts you print.

There are times when it's not feasible to use the editor for printing. For example, there is a limit on how big a file ed can handle (about 65,000 characters or 4000 lines). Secondly, it will only print one file at a time, and sometimes you want to print several, one after another. So here are a couple of alternatives.

First is cat, the simplest of all the printing programs. cat simply copies all the files in a list onto the terminal. So you can say

```
cat junk
```

or, to print two files,

```
cat junk temp
```

The two files are simply concatenated (hence the name "cat") onto the terminal.

pr produces formatted printouts of files. As with cat, pr prints all the files in a list. The difference is that it produces headings with date, time, page number and file name at the top of each page, and extra lines to skip over the fold in the paper. Thus,

```
pr junk temp
```

will list "junk" neatly, then skip to the top of a new page and list "temp" neatly.

pr will also produce multi-column output:

```
pr -3 junk
```

prints "junk" in 3-column format. You can use any reasonable number in place of "3" and **pr** will do its best.

It should be noted that **pr** is *not* a formatting program in the sense of shuffling lines around and justifying margins. The true formatters are **roff**, **nroff**, and **troff**, which we will get to in the section on document preparation.

There are also programs that print files on a high-speed printer. Look in your manual under **opr** and **lpr**. Which to use depends on the hardware configuration of your machine.

Shuffling Files About

Now that you have some files in the file system and some experience in printing them, you can try bigger things. For example, you can move a file from one place to another (which amounts to giving a file a new name), like this:

```
mv junk precious
```

This means that what used to be "junk" is now "precious". If you do an **ls** command now, you will get

```
precious
temp
```

Beware that if you move a file to another one that already exists, the already existing contents are lost forever.

If you want to make a *copy* of a file (that is, to have two versions of something), you can use the **cp** command:

```
cp precious temp1
```

makes a duplicate copy of "precious" in "temp1".

Finally, when you get tired of creating and moving files, there is a command to remove files from the file system, called **rm**.

```
rm temp temp1
```

will remove all of the files named. You will get a warning message if one of the named files wasn't there.

Filename, What's in a

So far we have used filenames without ever saying what's a legal name, so it's time for a couple of rules. First, filenames are limited to 14 characters, which is enough to be descriptive. Second, although you can use almost any charac-

ter in a filename, common sense says you should stick to ones that are visible, and that you should probably avoid characters that might be used with other meanings. We already saw, for example, that in the **ls** command, "**ls -l**" meant to list in time order. So if you had a file whose name was "-l", you would have a tough time listing it by name. There are a number of other characters which have special meaning either to **UNIX** as a whole or to numerous commands. To avoid pitfalls, you would probably do well to use only letters, numbers and the period. (Don't use the period as the first character of a filename, for reasons too complicated to go into.)

On to some more positive suggestions. Suppose you're typing a large document like a book. Logically this divides into many small pieces, like chapters and perhaps sections. Physically it must be divided too, for **ed** will not handle big files. Thus you should type the document as a number of files. You might have a separate file for each chapter, called

```
chap1
chap2
etc...
```

Or, if each chapter were broken into several files, you might have

```
chap1.1
chap1.2
chap1.3
...
chap2.1
chap2.2
...
```

You can now tell at a glance where a particular file fits into the whole.

There are advantages to a systematic naming convention which are not obvious to the novice **UNIX** user. What if you wanted to print the whole book? You could say

```
pr chap1.1 chap1.2 chap1.3 .....
```

but you would get tired pretty fast, and would probably even make mistakes. Fortunately, there is a shortcut. You can say

```
pr chap*
```

The "*" means "anything at all", so this translates into "print all files whose names begin with 'chap' ", listed in alphabetical order. This shorthand notation is not a property of the **pr** command, by the way. It is system-wide, a service of the program that interprets commands (the "shell" **sh(1)**). Using that fact, you can see

how to list the files of the book:

```
ls chap*
```

produces

```
chap1.1
chap1.2
chap1.3
...
```

The "*" is not limited to the last position in a filename — it can be anywhere. Thus

```
rm *junk*
```

removes all files that contain "junk" as any part of their name. As a special case, "*" by itself matches every filename, so

```
pr *
```

prints all the files (alphabetical order), and

```
rm *
```

removes *all files*. (You had better be sure that's what you wanted to say!)

The "*" is not the only pattern-matching feature available. Suppose you want to print only chapters 1 through 4 and 9 of the book. Then you can say

```
pr chap[12349]*
```

The "[...]" means to match any of the characters inside the brackets. You can also do this with

```
pr chap[1-49]*
```

"[a-z]" matches any character in the range a through z. There is also a "?" character, which matches any single character, so

```
pr ?
```

will print all files which have single-character names.

Of these niceties, "*" is probably the most useful, and you should get used to it. The others are frills, but worth knowing.

If you should ever have to turn off the special meaning of "*", "?", etc., enclose the entire argument in quotes (single or double), as in

```
ls "?"
```

What's in a Filename, Continued

When you first made that file called "junk", how did UNIX know that there wasn't another "junk" somewhere else, especially since the person in the next office is also reading this tutorial? The reason is that generally each user of UNIX has his own "directory", which contains

only the files that belong to him. When you create a new file, unless you take special action, the new file is made in your own directory, and is unrelated to any other file of the same name that might exist in someone else's directory.

The set of all files that UNIX knows about are organized into a (usually big) tree, with your files located several branches up into the tree. It is possible for you to "walk" around this tree, and to find any file in the system, by starting at the root of the tree and walking along the right set of branches.

To begin, type

```
ls /
```

"/" is the name of the root of the tree (a convention used by UNIX). You will get a response something like this:

```
bin
dev
etc
lib
tmp
usr
```

This is a collection of the basic directories of files that UNIX knows about. On most systems, "usr" is a directory that contains all the normal users of the system, like you. Now try

```
ls /usr
```

This should list a long series of names, among which is your own login name. Finally, try

```
ls /usr/your-name
```

You should get what you get from a plain

```
ls
```

Now try


```
cat /usr/your-name/junk
```

(if "junk" is still around). The name

```
/usr/your-name/junk
```

is called the "pathname" of the file that you normally think of as "junk". "Pathname" has an obvious meaning: it represents the full name of the path you have to follow through the tree of directories to get to a particular file. It is a universal rule in UNIX that anywhere you can use an ordinary filename, you can use a pathname.

Here is a picture which may make this clearer:



Notice that Mary's "junk" is unrelated to Eve's.

This isn't too exciting if all the files of interest are in your own directory, but if you work with someone else or on several projects concurrently, it becomes handy indeed. For example, your friends can print your book by saying

```
pr /usr/your-name/chap*
```

Similarly, you can find out what files your neighbor has by saying

```
ls /usr/neighbor-name
```

or make your own copy of one of his files by

```
cp /usr/your-neighbor/his-file yourfile
```

(If your neighbor doesn't want you poking around in his files, or vice versa, privacy can be arranged. Each file and directory can have read-write-execute permissions for the owner, a group, and everyone else, to control access. See **ls(1)** and **chmod(1)** for details. As a matter of observed fact, most users most of the time find openness of more benefit than privacy.)

As a final experiment with pathnames, try

```
ls /bin /usr/bin
```

Do some of the names look familiar? When you run a program, by typing its name after a "%", the system simply looks for a file of that name. It looks first in your directory (where it typically doesn't find it), then in "/bin" and finally in "/usr/bin". There is nothing magic about commands like **cat** or **ls**, except that they have been collected into two places to be easy to find and administer.

What if you work regularly with someone else on common information in his directory? You could just log in as your friend each time you want to, but you can also say "I want to work on his files instead of my own". This is done by changing the directory that you are currently in:

```
chdir /usr/your-friend
```

Now when you use a filename in something like **cat** or **pr**, it refers to the file in "your-friend's" directory. Changing directories doesn't affect any permissions associated with a file — if you couldn't access a file from your own directory, changing to another directory won't alter that fact.

If you forget what directory you're in, type **pwd**

("print working directory") to find out.

It is often convenient to arrange one's files so that all the files related to one thing are in a directory separate from other projects. For example, when you write your book, you might want to keep all the text in a directory called **book**. So make one with

```
mkdir book
```

then go to it with

```
chdir book
```

then start typing chapters. The book is now found in (presumably)

```
/usr/your-name/book
```

To delete a directory, see **rmdir(1)**.

You can go up one level in the tree of files by saying

```
chdir ..
```

".." is the name of the parent of whatever directory you are currently in. For completeness, "." is an alternate name for the directory you are in.

Using Files instead of the Terminal

Most of the commands we have seen so far produce output on the terminal; some, like the editor, also take their input from the terminal. It is universal in UNIX that the terminal can be replaced by a file for either or both of input and output. As one example, you could say

```
ls
```

to get a list of files. But you can also say

```
ls >filelist
```

to get a list of your files in the file "filelist". ("filelist" will be created if it doesn't already exist, or overwritten if it does.) The symbol ">" is used throughout UNIX to mean "put the output on the following file, rather than on the terminal". Nothing is produced on the terminal. As another example, you could concatenate several files into one by capturing the output of **cat** in a file:

```
cat f1 f2 f3 >temp
```

Similarly, the symbol "<" means to take the input for a program from the following file, instead of from the terminal. Thus, you could make up a script of commonly used editing commands and put them into a file called "script". Then you can run the script on a file by saying

```
ed file <script
```

Pipes

One of the novel contributions of UNIX is the idea of a *pipe*. A pipe is simply a way to connect the output of one program to the input of another program, so the two run as a sequence of processes — a pipe-line.

For example,

```
pr f g h
```

will print the files "f", "g" and "h", beginning each on a new page. Suppose you want them run together instead. You could say

```
cat f g h >temp
pr temp
rm temp
```

but this is more work than necessary. Clearly what we want is to take the output of `cat` and connect it to the input of `pr`. So let us use a pipe:

```
cat f g h | pr
```

The vertical bar means to take the output from `cat`, which would normally have gone to the terminal, and put it into `pr`, which formats it neatly.

Any program that reads from the terminal can read from a pipe instead; any program that writes on the terminal can drive a pipe. You can have as many elements in a pipeline as you wish.

Many UNIX programs are written so that they will take their input from one or more files if file arguments are given; if no arguments are given they will read from the terminal, and thus can be used in pipelines.

The Shell

We have already mentioned once or twice the mysterious "shell," which is in fact `sh(1)`. The shell is the program that interprets what you type as commands and arguments. It also looks after translating "*", etc., into lists of filenames.

The shell has other capabilities too. For example, you can start two programs with one command line by separating the commands with

a semicolon; the shell recognizes the semicolon and breaks the line into two commands. Thus

```
date; who
```

does both commands before returning with a "%".

You can also have more than one program running *simultaneously* if you wish. For example, if you are doing something time-consuming, like the editor script of an earlier section, and you don't want to wait around for the results before starting something else, you can say

```
ed file <script &
```

The ampersand at the end of a command line says "start this command running, then take further commands from the terminal immediately." Thus the script will begin, but you can do something else at the same time. Of course, to keep the output from interfering with what you're doing on the terminal, it would be better to have said

```
ed file <script >lines &
```

which would save the output lines in a file called "lines"

When you initiate a command with "&", UNIX replies with a number called the process number, which identifies the command in case you later want to stop it. If you do, you can say

```
kill process-number
```

You might also read `ps(1)`.

You can say

```
(command-1; command-2; command-3) &
```

to start these commands in the background, or you can start a background pipeline with

```
command-1 | command-2 &
```

Just as you can tell the editor or some similar program to take its input from a file instead of from the terminal, you can tell the shell to read a file to get commands. (Why not? The shell after all is just a program, albeit a clever one.) For instance, suppose you want to set tabs on your terminal, and find out the date and who's on the system every time you log in. Then you can put the three necessary commands (`tabs`; `date`; `who`) into a file, let's call it "xxx", and then run it with either

```
sh xxx
```

or

```
sh <xxx
```


This says to run the shell with the file "xxx" as input. The effect is as if you had typed the contents of "xxx" on the terminal. (If this is to be a regular thing, you can eliminate the need to type "sh"; see **chmod** (1) and **sh** (1).)

The shell has quite a few other capabilities as well, some of which we'll get to in the section on programming.

III. DOCUMENT PREPARATION

UNIX is extensively used for document preparation. There are three major *formatting* programs, that is, programs which produce a text with justified right margins, automatic page numbering and titling, automatic hyphenation, and the like. The simplest of these formatters is **roff**, which in fact is simple enough that if you type almost any text into a file and "roff" it, you will get plausibly formatted output. You can do better with a little knowledge, but basically it's easy to learn and use. We'll get back to **roff** shortly.

nroff is similar to **roff** but does much less for you automatically. It will do a great deal more, once you know how to use it.

Both **roff** and **nroff** are designed to produce output on terminals, line-printers, and the like. The third formatter, **troff** (pronounced "tee-roff"), instead drives a Graphic Systems phototypesetter, which produces very high quality output on photographic paper. This paper was printed on the phototypesetter by **troff**.

Because **nroff** and **troff** are relatively hard to learn to use effectively, several "packages" of canned formatting requests are available which let you do things like paragraphs, running titles, multi-column output, and so on, with little effort. Regrettably, details vary from system to system.

ROFF

The basic idea of **roff** (and of **nroff** and **troff**, for that matter) is that the text to be formatted contains within it "formatting commands" that indicate in detail how the formatted text is to look. For example, there might be commands that specify how long lines are, whether to use single or double spacing, and what running titles to use on each page. In general, you don't have to spell out all of the possible formatting details. Most of them have "default values", which you will get if you say nothing at all. For example, unless you take special precautions, you'll get single-spaced output, 65-character lines, justified right margins, and 58

text lines per page when you **roff** a file. This is the reason that **roff** is so simple — most of the decisions have already been made for you.

Some things do have to be done, however. If you want a document broken into paragraphs, you have to tell **roff** where to add the extra blank lines. This is done with the ".sp" command:

```
this is the end of one paragraph.
.sp
This begins the next paragraph ...
```

In **roff** (and in **nroff** and **troff**), formatting commands consist of a period followed by two letters, and they must appear at the beginning of a line, all by themselves. The ".sp" command tells **roff** to finish printing any of the previous line that might be still unprinted, then print a blank line before continuing. You can have more space if you wish; ".sp 2" asks for 2 spaces, and so on.

If you simply want to ensure that subsequent text appears on a fresh output line, you can use the command ".br" (for "break") instead of ".sp".

Most of the other commonly-used **roff** commands are equally simple. For example you can center one or more lines with the ".ce" command.

```
.ce
Title of Paper
.sp 2
```

causes the title to be centered, then followed by two blank lines. As with ".sp", ".ce" can be followed by a number; in that case, that many input lines are centered.

".ul" underlines lines, and can also be followed by a number:

```
.ce 2
.ul 2
An Earth-shaking Paper
.sp
John O Scientist
```

will center and underline the two text lines. Notice that the ".sp" between them is not part of the line count.

You can get multiple-line spacing instead of the default single-spacing with the ".ls" command:

```
.ls 2
```

causes double spacing.

If you're typing things like tables, you will not want the automatic filling-up and justification of output lines that is done by default. You can turn this off with the command ".nf" (no-fill), and then back on again with ".fi" (fill). Thus

```
.nf
here lines will appear just
as you typed them —
no extra spaces, no moving of words.
.fi
Now go back to filling up output lines.
```

You can change the line-length with ".ll", and the left margin (the indent) by ".in". These are often used together to make offset blocks of text:

```
.ll -10
.in +10
this text will be moved 10
spaces to the right and the
lines will also be shortened 10
characters from the right. The
"+" and "-" mean to change
the previous value by that
much. Now revert:
.ll +10
.in -10
```

Notice that ".ll +10" adds ten characters to the line length, while ".ll 10" makes the line ten characters *long*.

The ".ti" command indents (in either direction) just like ".in", except for only one line. Thus to make a new paragraph with a 10-character indent, you would say

```
.sp
.ti +10
New paragraph ...
```

You can put running titles on both top and bottom of each page, like this:

```
.he "left top"center top"right top"
.fo "left bottom"center bottom"right bottom"
```

The header or footer is divided into three parts, which are marked off by any character you like. (We used a double quote.) If there's nothing between the markers, that part of the title will be blank. If you use a percent sign anywhere in ".he" or ".fo", the current page number will be inserted. So to get centered page numbers with dashes around them, at the top, use

```
.he ""- % -""
```

You can skip to the top of a new page at any time with the ".bp" command; if ".bp" is followed by a number, that will be the new page number.

The foregoing is probably enough about **roff** for you to go off and format most everyday documents. Read **roff** (1) for more details.

Hints for Preparing Documents

Most documents go through several versions (always more than you expected) before they are finally finished. Accordingly, you should do whatever possible to make the job of changing them easy.

First, when you do the purely mechanical operations of typing, type so subsequent editing will be easy. Start each sentence on a new line. Make lines short, and break lines at natural places, such as after commas and semicolons, rather than randomly. Since most people change documents by rewriting phrases and adding, deleting and rearranging sentences, these precautions simplify any editing you have to do later.

The second aspect of making change easy is not to commit yourself to formatting details too early. For example, if you decide that each paragraph is to have a space and an indent of 10 characters, you might type, before each,

```
.sp
.ti +10
```

But what happens when later you decide that it would have been better to have no space and an indent of only 5 characters? It's tedious indeed to go back and patch this up.

Fortunately, all of the formatters let you delay decisions until the actual moment of running. The secret is to define a new operation (called a *macro*), for each formatting operation you want to do, like making a new paragraph. You can say, in all three formatters,

```
.de PP
.sp
.ti +10
..
```

This *defines* ".PP" as a new **roff** (or **nroff** or **troff**) operation, whose meaning is exactly

```
.sp
.ti +10
```

(The "." marks the end of the definition.) Whenever ".PP" is encountered in the text, it is as if you had typed the two lines of the definition in place of it.

The beauty of this scheme is that now, if you change your mind about what a paragraph should look like, you can change the formatted output merely by changing the definition of "PP" and re-running the formatter.

As a rule of thumb, for all but the most trivial jobs, you should type a document in terms of a set of macros like "PP", and then define them appropriately. As long as you have entered the text in some systematic way, it can always be cleaned up and re-formatted by a judicious combination of editing and macro definitions. The packages of formatting commands that we mentioned earlier are simply collections of macros designed for particular formatting tasks.

One of the main differences between **roff** and the other formatters is that macros in **roff** can only be lines of text and formatting commands. In **nroff** and **troff**, macros may have arguments, so they can have different effects depending on how they are called (in exactly the same way that the ".sp" command has an argument, the number of spaces you want).

Miscellany

In addition to the basic formatters, UNIX provides a host of supporting programs. **eqn** and **neqn** let you integrate mathematics into the text of a document, in a language that closely resembles the way you would speak it aloud. **spell** and **typo** detect possible spelling mistakes in a document. **grep** looks for lines containing a particular text pattern (rather like the editor's context search does, but on a whole series of files). For example,

```
grep "ing$" chap*
```

will find all lines ending in the letters "ing" in the series of files "chap*". (It is almost always a good practice to put quotes around the pattern you're searching for, in case it contains characters that have a special meaning for the shell.)

wc counts the words and (optionally) lines in a set of files. **tr** translates characters into other characters; for example it will convert upper to lower case and vice versa. This translates upper into lower:

```
tr "[A-Z]" "[a-z]"
```

diff prints a list of the differences between two files, so you can compare two versions of something automatically (which certainly beats proofreading by hand). **sort** sorts files in a variety of ways; **cref** makes cross-references; **ptx** makes a permuted index (keyword-in-context listing).

Most of these programs are either independently documented (like **eqn** and **neqn**), or are sufficiently simple that the description in the *UNIX Programmer's Manual* is adequate explanation.

IV. PROGRAMMING

UNIX is a marvelously pleasant and productive system for writing programs; productivity seems to be an order of magnitude higher than on other interactive systems.

There will be no attempt made to teach any of the programming languages available on UNIX, but a few words of advice are in order. First, UNIX is written in C, as is most of the applications code. If you are undertaking anything substantial, C is the only reasonable choice. More on that in a moment. But remember that there are quite a few programs already written, some of which have substantial power.

The editor can be made to do things that would normally require special programs on other systems. For example, to list the first and last lines of each of a set of files, say a book, you could laboriously type

```
ed
e chap1.1
lp
Sp
e chap1.2
lp
Sp
etc.
```

But instead you can do the job once and for all. Type

```
ls chap* >temp
```

to get the list of filenames into a file. Then edit this file to make the necessary series of editing commands (using the global commands of **ed**), and write it into "script". Now the command

```
ed <script
```

will produce the same output as the laborious hand typing.

The pipe mechanism lets you fabricate quite complicated operations out of spare parts already built. For example, the first draft of the **spell** program was (roughly)

```

cat ... (collect the files)
| tr ... (put each word on a new line,
         delete punctuation, etc.)
| sort (into dictionary order)
| uniq (strip out duplicates)
| comm (list words found in text but
        not in dictionary)

```

Programming the Shell

An option often overlooked by newcomers is that the shell is itself a programming language, and since UNIX already has a host of building-block programs, you can sometimes avoid writing a special purpose program merely by piecing together some of the building blocks with shell command files.

As an unlikely example, suppose you want to count the number of users on the machine every hour. You could type

```

date
who | wc -l

```

every hour, and write down the numbers, but that is rather primitive. The next step is probably to say

```
(date; who | wc -l) >>users
```

which uses ">>" to *append* to the end of the file "users". (We haven't mentioned ">>" before — it's another service of the shell.) Now all you have to do is to put a loop around this, and ensure that it's done every hour. Thus, place the following commands into a file, say "count":

```

: loop
(date; who | wc -l) >>users
sleep 3600
goto loop

```

The command `:` is followed by a space and a label, which you can then `goto`. Notice that it's quite legal to branch backwards. Now if you issue the command

```
sh count &
```

the users will be counted every hour, and you can go on with other things. (You will have to use `kill` to stop counting.)

If you would like "every hour" to be a parameter, you can arrange for that too:

```

: loop
(date; who | wc -l) >>users
sleep $1
goto loop

```

"\$1" means the first argument when this procedure is invoked. If you say

```
sh count 60
```

it will count every minute. A shell program can have up to nine arguments, "\$1" through "\$9".

The other aspect of programming is conditional testing. The `if` command can test conditions and execute commands accordingly. As a simple example, suppose you want to add to your login sequence something to print your mail if you have some. Thus, knowing that mail is stored in a file called 'mailbox', you could say

```
if -r mailbox mail
```

This says "if the file 'mailbox' is readable, execute the `mail` command."

As another example, you could arrange that the "count" procedure count every hour by default, but allow an optional argument to specify a different time. Simply replace the "sleep \$1" line by

```
if $1x = x sleep 3600
if $1x != x sleep $1
```

The construction

```
if $1x = x
```

tests whether "\$1", the first argument, was present or absent.

More complicated conditions can be tested: you can find out the status of an executed command, and you can combine conditions with 'and', 'or', 'not' and parentheses — see `if(1)`. You should also read `shift(1)` which describes how to manipulate arguments to shell command files.

Programming in C

As we said, C is the language of choice: everything in UNIX is tuned to it. It is also a remarkably easy language to use once you get started. Sections II and III of the manual describe the system interfaces, that is, how you do I/O and similar functions.

You can write quite significant C programs with the level of I/O and system interface described in *Programming in C: A Tutorial*, if you use existing programs and pipes to help. For example, rather than learning how to open and close files you can (at least temporarily) write a program that reads from its standard input, and use `cat` to concatenate several files into it. This may not be adequate for the long run, but for the early stages it's just right.

There are a number of supporting programs that go with C. The C debugger, `cdb`, is marginally useful for digging through the dead

bodies of C programs. **db**, the assembly language debugger, is actually more useful most of the time, but you have to know more about the machine and system to use it well. The most effective debugging tool is still careful thought, coupled with judiciously placed print statements.

You can instrument C programs and thus find out where they spend their time and what parts are worth optimising. Compile the routines with the "-p" option; after the test run use **prof** to print an execution profile. The command **time** will give you the gross run-time statistics of a program, but it's not super accurate or reproducible.

C programs that don't depend too much on special features of UNIX can be moved to the Honeywell 6070 and IBM 370 systems with modest effort. Read *The GCOS C Library* by M. E. Lesk and B. A. Barres for details.

Miscellany

If you *have* to use Fortran, you might consider **ratfor**, which gives you the decent control structures and free-form input that characterize C, yet lets you write code that is still portable to other environments. Bear in mind that UNIX Fortran tends to produce large and relatively slow-running programs. Furthermore, supporting software like **db**, **prof**, etc., are all virtually useless with Fortran programs.

If you want to use assembly language (all heavens forfend!), try the implementation language LIL, which gives you many of the advantages of a high-level language, like decent control flow structures, but still lets you get close to the machine if you really want to.

If your application requires you to translate a language into a set of actions or another language, you are in effect building a compiler, though probably a small one. In that case, you should be using the **yacc** compiler-compiler, which helps you develop a compiler quickly.

V. UNIX READING LIST

General:

UNIX Programmer's Manual (Ken Thompson, Dennis Ritchie, and a cast of thousands). Lists commands, system routines and interfaces, file formats, and some of the maintenance procedures. You can't live without this, although you will probably only read section I.

The UNIX Time-sharing System (Ken Thompson, Dennis Ritchie). CACM, July 1974. An overview of the system, for people interested in operating systems. Worth reading by anyone

who programs. Contains a remarkable number of one-sentence observations on how to do things right.

Document Preparation:

A Tutorial Introduction to the UNIX Text Editor. (Brian Kernighan). Bell Laboratories internal memorandum. Weak on the more esoteric uses of the editor, but still probably the easiest way to learn **ed**.

Typing Documents on UNIX. (Mike Lesk). Bell Laboratories internal memorandum. A macro package to isolate the novice from the vagaries of the formatting programs. If this specific package isn't available on your system, something similar probably is. This one works with both **nroff** and **troff**.

Programming:

Programming in C: A Tutorial (Brian Kernighan). Bell Laboratories internal memorandum. The easiest way to start learning C, but it's no help at all with the interface to the system beyond the simplest IO. Should be read in conjunction with

C Reference Manual (Dennis Ritchie). Bell Laboratories internal memorandum. An excellent reference, but a bit heavy going for the beginner, especially one who has never used a language like C.

Others:

D. M. Ritchie, UNIX Assembler Reference Manual.

B. W. Kernighan and L. L. Cherry, A System for Typesetting Mathematics, Computing Science Tech. Rep. 17.

M. E. Lesk and B. A. Barres, The GCOS C Library. Bell Laboratories internal memorandum.

K. Thompson and D. M. Ritchie, Setting Up UNIX.

M. D. McIlroy, UNIX Summary.

D. M. Ritchie, The UNIX I/O System.

A. D. Hall, The M6 Macro Processor, Computing Science Tech. Rep. 2.

J. F. Ossanna, NROFF User's Manual — Second Edition, Bell Laboratories internal memorandum.

D. M. Ritchie and K. Thompson, Regenerating System Software.

B. W. Kernighan, Ratfor—A Rational Fortran, Bell Laboratories internal memorandum.

M. D. McIlroy, Synthetic English Speech by Rule, Computing Science Tech. Rep. 14.

M. D. McIlroy, A Manual for the TMG Compiler-writing Language. Bell Laboratories

internal memorandum.

J. F. Ossanna, TROFF Users' Manual, Bell Laboratories internal memorandum.

B. W. Kernighan, TROFF Made Trivial, Bell Laboratories internal memorandum.

R. H. Morris and L. L. Cherry, Computer Detection of Typographical Errors, Computing Science Tech. Rep. 18.


S. C. Johnson, YACC (Yet Another Compiler-Compiler), Bell Laboratories internal memorandum.


P. J. Plauger, Programming in LIL: A Tutorial, Bell Laboratories internal memorandum.

Index

& (asynchronous process) 8
; (multiple processes) 8
* (pattern match) 5
[] (pattern match) 6
? (pattern match) 6
<> (redirect I/O) 7
>> (file append) 12
backslash (\) 2
cat (concatenate files) 4
cdb (C debugger) 12
chdir (change directory) 7
chmod (change protection) 7
command arguments 4
command files 8
cp (copy files) 5
cref (cross reference) 11
date 2
db (assembly debugger) 13
delete (DEL) 2
diff (file comparison) 11
directories 7
document formatting 9
ed (editor) 3
editor programming 11
EOT (end of file) 3
eqn (mathematics) 11
erase character (#) 2
file system structure 6
filenames 5
file protection 7
goto 12
grep (pattern matching) 11
if (condition test) 12
index 14
kill a program 8
kill a character (@) 2

lil (high-level assembler) 13
login 1
logout 2
ls (list file names) 4
macro for formatting 10
mail 2
multi-columns printing (pr) 5
mv (move files) 5
nroff 9
on-line manual 3
opr (offline print) 5
pathname 6
pattern match in filenames 5
pipes (|) 8
pr (print files) 4
prof (run-time monitor) 13
protection 7
ptx (permuted index) 11
pwd (working directory) 7
quotes 6
ratfor (decent Fortran) 13
readahead 2
reading list 13
redirect I/O (<>) 7
RETURN key 1
rm (remove files) 5
rmdir (remove directory) 7
roff (text formatting) 9
root (of file system) 6
shell (command interpreter) 8
shell arguments (\$) 12
shell programming 12
shift (shell arguments) 12
sleep 12
sort 11
spell (find spelling mistakes)
stopping a program 2
stty (set terminal options) 2
tabs (set tab stops) 2
terminal types 1
time (time programs) 13
tr (translate characters) 11
troff (typesetting) 9
typo (find spelling mistakes) 11
wc (word count) 11
who (who is logged in) 2
write (to a user) 3
yacc (compiler-compiler) 13





UNIX Programming

Brian W. Kernighan

Dennis M. Ritchie

Bell Laboratories,
Murray Hill, New Jersey 07974

ABSTRACT

This paper is an introduction to programming on UNIX. The emphasis is on how to write programs that interface to the operating system. The topics discussed include

- handling command arguments
- rudimentary I/O; the standard input and output
- the portable C I/O library; file system access
- standard UNIX file I/O
- low-level I/O
- executing commands from programs
- signals — interrupts, etc.



UNIX Programming

Brian W. Kernighan

Dennis M. Ritchie

Bell Laboratories,
Murray Hill, New Jersey 07974

1. Introduction

This paper describes how to write UNIX programs that interface with the operating system in a non-trivial way. This includes programs that use files by name, that do large amounts of input or output, that invoke other programs as they run, or that attempt to catch interrupts and other signals during execution.

The document collects material which is scattered throughout several sections of *The UNIX Programmer's Manual* [1]. There is no attempt to be complete; only generally useful material is dealt with. It is assumed that you will be programming in C, so you will have to be able to read the language roughly up to the level of *Programming in C — A Tutorial* [2]. You should also be familiar with UNIX itself to the level of *UNIX for Beginners* [3].

2. Program Arguments

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

Here is a program that simply echoes its arguments back to the terminal. (This is much like the command `echo`.)

```
main( argc, argv )
int argc;
char *argv[ ];
{
    int i;
    for( i=1; i < argc; i++ )
        printf( "%s ", argv[i] );
    putchar( '\n' );
}
```

`main` is called with two arguments, the argument count and the array of arguments. `argv` is a pointer to an array, whose individual elements are pointers to arrays of characters; each is terminated by `'\0'`, so they can be treated as strings. `argv[0]` is the name of the command itself, so we start by printing `argv[1]` and loop until we've printed them all. Each `argv[i]` is a character array, so we use a `'%s'` in the `printf`.

A common convention in UNIX programs is that an argument which begins with `'-'` indicates a flag or option of some sort. For example, suppose we want a program to be callable as

```
prog -abc arg1 arg2 ...
```

where the '-' argument is optional; if it is present, it must be first and may be followed by any combination of the options 'a', 'b', and 'c'.

```
main(argc, argv)
int argc;
char *argv[ ];
{
    ...
    aflag = bflag = cflag = 0;
    while( argc > 1 && argv[1][0] == '-' ) {
        for( i=1; (c=argv[1][i]) != '\0'; i++ )
            if( c=='a' )
                aflag++;
            else if( c=='b' )
                bflag++;
            else if( c=='c' )
                cflag++;
            else
                printf( "%c?\n", c );
        --argc;
        ++argv;
    }
    ...
}
```

The statements

```
--argc;
++argv;
```

drop the first argument from the list and adjust the count, so after interpreting the flag argument, the rest of the program is independent of whether or not it existed. This works because **argv** is a pointer which can be incremented. Notice also that, for greatest generality, the **while** and **for** loops in combination allow the user to write the options either as

```
program -a -b -c
```

or as

```
program -abc
```

This degree of generality is unfortunately not very common; most commands require one form or the other.

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

3. Rudimentary Input and Output

The next several sections will discuss various aspects of input/output, including how to create, open, and close files from programs. There are several ways to do most of these operations; these sections are organized so easy things are described first.

3.1. The "Standard Input" and "Standard Output"

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

```
prog <file
```

causes `prog` to read `file` instead of the terminal. `prog` itself knows nothing about where its input is coming from. This is also true if the input comes from another program via the UNIX pipe mechanism:

```
otherprog | prog
```

will provide the input for `prog` from the output of `otherprog`.

`getchar` returns the value zero (often written as the null-character '\0') when it encounters the end of file (or an error) on whatever you are reading. Bear in mind that '\0' may be a legitimate value in some contexts. If it is, you can't use `getchar`; read ahead to the discussion of `getc`.

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

```
prog >outfile
```

will write the output onto `outfile` instead of the terminal. And a pipe can be used:

```
prog | otherprog
```

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

The function `printf`, which formats output in various ways, uses `putchar` to finally print the output, so output produced by `printf` also finds its way to the standard output.

A surprising number of programs read only one input and write one output; for such programs I/O with `getchar`, `putchar`, and `printf` may be entirely adequate, and it is almost always enough to get started. This is particularly true, given the UNIX pipe facility for connecting the output of one program to the input of the next. For example, here is a complete program that acts as a "filter" to strip out all ascii control characters from its input (except for newline and tab).

```
main( ) {
    int c;
    while( c=getchar( ) )
        if( (c>=' ' && c<0177) || c=='\t' || c=='\n' )
            putchar(c);
    exit( 0 );
}
```

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

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

and thus avoid learning how to access files from a program. By the way, the call to `exit` at the end is not necessary to make the program work properly, but it assures that any caller of the program will see a normal termination status (conventionally 0) from the program when it completes. Section 7.3 discusses status returns in more detail.

3.2. File Descriptors

Before we go much further into our description of I/O, we have to talk about file descriptors. Any program which does any input or output does so by reading or writing files. This is true even though the file in question may actually be a device like the user's terminal. Associated with each file being used for input or output is a small non-negative integer called a "file descriptor;" whenever I/O is to be done on the file, the file descriptor is used to identify the file, instead of the name. (This is roughly analogous to the use of `READ(5,...)` and `WRITE(6,...)` in Fortran.)

In the most general case, to do I/O on a file you have to

Open the file for reading and/or writing — UNIX connects the name of the file with a file descriptor which it generates and returns to you. An alternative form of *open* will create the file if it doesn't exist.

Read or write on the file.

Close the file, which breaks the connection between name and descriptor.

And finally, you may want to

Buffer input and output if necessary for efficiency.

In the simplest cases, like **getchar** and **putchar**, all opening and closing is done for you; you only have to worry about reading and writing the right things. In more complicated situations, you have to do more work, but you gain flexibility.

getchar and **putchar** depend on the fact that when the command interpreter (the shell) runs a program, it opens three files, with file descriptors 0, 1, and 2. All of these are normally connected to the terminal, so if you read 0 and write 1 or 2, I/O is done on the terminal. If you use '<' or '>', the shell changes the default assignment for file 0 or 1 from your terminal to the named file, and opens it for you. This way your program need not know where its input comes from nor where its output goes, so long as it uses files 0 and 1. Naturally, **getchar** reads 0 and **putchar** writes 1. When the program terminates, the files are closed automatically. (We'll get back to file 2 in Section 3.4.)

3.3. Buffering the Standard Input and Output

If you are producing large amounts of output, you may find that programs which use **putchar** are slow, because each call to it actually requires a system call. You can speed up such a program markedly by buffering the output. (Generally UNIX buffers input automatically, which is almost always what is wanted. See Section 6 for how to turn off buffering.)

Buffering requires a little witchcraft. First we need a buffer area. This can be provided by referring to an external variable **fout** which is used by **putchar** for output. **fout** is actually a structure defined as part of **putchar**, but its internal structure needn't concern us now. Second, we have to connect the standard output to **fout**. Lastly, when all processing is done, any output remaining in the output buffer has to be flushed out; this is not done automatically. Putting all this together gives

```
main( ) {
    extern int fout;
    ...
    fout = dup( 1 );      /* buffer standard output */
    [ processing ... ]
    flush( ); /* force out last buffer */
    exit( 0 );
}
```

This should probably be treated as black magic for now. Briefly, `putchar` buffers only if the file has a descriptor greater than 2. The call to `dup` creates a new file descriptor that refers to the same file as the original 1, but is guaranteed to be buffered. `flush` forces out everything that has collected in the buffer.

As an aside, there is an external variable `fin` used by `getchar` much as `putchar` uses `fout`.

3.4. Diagnostic Output — The Error File

If we are going to run our control-character stripping program, it might be nice to know whether it actually removed any characters. But at the same time we can't just print a message, because that will be sent to the same place as the data itself. That is, if we say

```
ccstrip <infile >outfile
```

we don't want `outfile` to contain a line saying "there were 3 bad characters."

Here is how to get this kind of diagnostic information separated from the standard output and placed on the terminal regardless. Just as the file descriptors 0 and 1 are predefined, so is file descriptor 2. Unless you go out of your way to change it, output written on file 2 will find its way to your terminal. Thus in simple cases, you can simply set `fout` to 2 to direct output to the terminal.

```
main( ) {
    extern int fout;
    ...
    flush( ); /* clear out standard output */
    fout = 2;
    printf( "%d bad characters\n", badchar );
    exit( 0 );
}
```

This can get clumsy if the diagnostic output is to show up during the running of the program instead of all at the end, because you have to flush output for one file before switching to another, *each* time you switch.

4. The Portable C Library

The portable C library [3] was written by Mike Lesk to provide a set of high-level I/O routines that could be implemented on any machine with a C compiler, and thus permit some degree of program transferability. C programs which use this library for I/O can be moved, with essentially no change, between UNIX, GCOS, and IBM-TSO. At the same time, many of the details of buffering, file access, etc., are hidden from the user. The routines are somewhat bigger and slower than the analogous routines that are a standard part of UNIX, but they do handle a number of things automatically which other packages do not.

If you use programs from the portable C library, you have to ask specifically that it be searched when you compile or link-edit your program by writing `-lp` at the end of the arguments to the `cc` command:

```
cc prog.c -lp
```

4.1. File Opening and Closing

Before a file can be read or written, it has to be opened with the routine `copen`. `copen` has two arguments, the file name, and the type of access wanted (read, write, or append). Thus:

```
fd = copen( "/usr/bwk/foo", 'r' );
```

opens `/usr/bwk/foo` for reading. The value returned by `copen` is the file descriptor assigned by UNIX, to be used later for reading the file. If this number is `-1`, an error has occurred, so some defensive action has to take place. The usual code is

```
if( (fd=copen(name,mode)) == -1 )
    error( "Can't open file", name );
```

where `error` is some message printer.

`mode` is one of `'r'`, `'w'`, or `'a'`. The arguments `'w'` and `'a'` mean writing and appending. If the file is to be opened for writing, and if it doesn't already exist, it will be created for you. If the argument is `'w'` and if the file already exists, it will be truncated to zero length. If the argument is `'a'`, you will write at the end of the file in either case. (And of course if any of this fails, you will get a `-1` error return.)

When you finish processing a file, you should close it explicitly with

```
cclose(fd);
```

where `fd` is the file descriptor handed you by `copen`. `cclose` will flush out any buffer contents remaining before closing the file; there is no `flush` command in the portable C library. Finally, when your program is done, you should call `cexit` which will close any open files and flush out their buffers; like `exit`, it then terminates the program and delivers its argument as termination status. By the way, there is a limit of 15 simultaneously open files, so if your program deals with many files it will have to call `cclose` explicitly.

Let us illustrate these routines with a program modeled after `grep`. `grep`, the general purpose pattern finder, has several arguments, some of which are file names. If there are no file names, it reads the standard input. And it writes on the standard output. That is,

```
grep pattern [optional list of input files]
```

will print all input lines that contain "pattern". Thus we write

```
main( argc, argv )
int argc;
char *argv[ ];
{
    char line[1000], *pattern;
    int i, fd;
    if( argc < 2 ) {
        printf( 2, "Usage: grep pattern [file..]\n" );
        cexit( 1 );
    }
    pattern = argv[1];
    i = 2;
    do {
        if( argc == 2 )
            fd = 0; /* use standard input */
        else if( (fd=copen(argv[i], 'r')) == -1 ){
            printf( 2, "can't open %s\n", argv[i] );
            cexit( 1 );
        }
    }
```



```
while( gets( line, fd ) )
    if( match( line, pattern ) )
        puts( line );
if( fd != 0 )
    cclose( fd );
} while( ++i < argc );
cexit( 0 );
}
```

First the arguments are validated. Then the files are gone through in order; each is opened (if possible), and scanned in the **while** loop.

gets and **puts** are routines in the portable C library that read and write a line at a time. **match** is an unspecified routine (you write it!) that tells whether the line contains the pattern.

Notice that we have written a call to **printf** with a first argument of 2. If the first argument of **printf** is a small positive integer, this is assumed to be a file descriptor, and the output is sent there. (Caveat: this is true in the portable library only, unfortunately.) Thus any error messages go to the user's terminal.

There are a couple of other things to note in passing. First, the basic design of the program is that it allows input to be either from a set of files or from the standard input, and it writes on the standard output. This way the program can be used stand-alone or as part of a pipeline. It is important to design and implement programs this way whenever possible. Second, the program signals errors in two ways. The diagnostic output goes out onto the error file so it finds its way to your terminal instead of disappearing down a pipeline or into a file. Also the program returns a value when it calls **[c]exit**, so the success or failure of the command can be tested from within another program that uses this one as a sub-process.

4.2. Character I/O

The portable library provides two routines for I/O of individual characters. **cgetc** and **cputc** are quite analogous to **getchar** and **putchar**, except that they require an additional argument to specify a file descriptor. Thus

```
cgetc( 0 ); cputc( c, 1 );
```

reads the standard input and writes the standard output, while

```
cgetc( fd ); cputc( c, fd );
```

reads and writes **fd**. For convenience, **getchar** and **putchar** are also provided; they just call **cgetc** or **cputc** with the appropriate file descriptor argument. As an exercise you might write your own versions of **gets** and **puts** using **cgetc** and **cputc**.

4.3. Miscellaneous

The portable C library contains several other goodies, the most useful of which is probably the function **scanf**, which provides input format conversion similar to **printf** on output; it will convert strings to integers, floating point numbers, and so on. There is also a way to use **printf** to do in-core format conversion:

```
printf( - 1, s, format, ... );
```

will put its output in string **s** instead of a file. The **ungetc** function can be used to push characters back onto the input stream for re-reading. And the **feof** function can be used to test explicitly for end-of-file, so data which contains null characters can be handled with **cgetc**.

5. Standard UNIX I/O

This section describes the I/O routines provided as part of "standard" UNIX. They do somewhat less for you than the portable library, but are more efficient, and perhaps more widely available. The essential difference is that the user has to supply the buffer for each file explicitly, and provide his own flushing of output before closing a file. These routines are described in sections II and III of the *UNIX Programmer's Manual* [1].

Let us illustrate by writing a simplified version of `cmp`, a program that compares two files byte by byte.

```
main( argc, argv )
int argc;
char *argv[ ];
{
    int c1, c2, byte;
    int buf1[259], buf2[259];
    if( argc != 3 )
        error( "Usage: cmp file1 file2" );
    if( fopen( argv[1], buf1 ) < 0 )
        error( "can't open %s", argv[1] );
    if( fopen( argv[2], buf2 ) < 0 )
        error( "can't open %s", argv[2] );
    for( byte=0 ; ; byte++ ) {
        c1 = getc( buf1 );
        c2 = getc( buf2 );
        if( c1 < 0 || c2 < 0 )
            break;
        if( c1 != c2 )
            printf( "%6u %3o %3o\n", byte, c1, c2 );
    }
    if( c1 == c2 )
        exit( 0 );
    if( c1 < 0 )
        printf( "eof 1\n" );
    else
        printf( "eof 2\n" );
    exit( 1 );
}

error( s1, s2 )
char *s1;
int s2;
{
    printf( s1, s2 );
    printf( "\n" );
    exit( 2 );
}
```

Files are opened with `fopen`, whose two arguments are the filename and a buffer, which is always declared `int[259]`. The buffer is actually used as a structure by the I/O routines (for example, the file descriptor goes into `buf[0]`), but we need not be concerned with that here; it is sufficient to note that the buffer is the connection between all the routines concerned with a file. As with `copen`, `fopen` returns a `-1` if the access failed for any reason. `error` is a simple routine to print out a message and exit with the appropriate status return.

`getc` reads the input using the `buffer` argument to tell it what file to read. On end of file, `getc` returns the value `-1`, not zero, so it can be used to read data containing explicit zero characters, where `getchar` and `cgetc` are not suitable.

The situation for output is slightly more complicated. First, the output file may not exist. The routine `fcreat` allows for this:

```
fcreat( name, buf )
```

creates the file if it doesn't exist; if it does exist, it is truncated to zero length. To write on the file,

```
putc( ch, buf )
```

puts the character `ch` onto the file. When writing is done, call

```
fflush( buf )
```

to force out any left-over output, then call

```
close( buf[0] )
```

to close the file. Termination of the program, by calling `exit` or otherwise, closes all open files, but *it does not flush the output buffers*.

As an aside, the way to delete a file is simply to call

```
unlink( filename );
```

This returns `-1` if the unlinking failed. Renaming a file uses both `link` and `unlink`:

```
if ( link( oldname, newname ) >= 0 )
    unlink( oldname );
```

renames the file from "oldname" to "newname," taking care not to delete the file if the link failed, as it would, for example, if "newname" already existed.

6. Low-Level I/O

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

6.1. I/O Proper

For the low-level I/O routines, files are identified directly by their file descriptors. Here is a simple version of `cp`, a program which copies one file to another.

```
main( argc, argv )
int argc;
char *argv[ ];
{
    int f1, f2, n;
    char buf[512];
    if( argc != 3 )
        error( "Usage: cp from to" );
    if( (f1=open(argv[1], 0) < 0 )
        error( "can't open %s", argv[1] );
    if( (f2=creat(argv[2], 0666)) < 0 )
        error( "can't open %s", argv[2] );
    while( (n=read(f1, buf, 512)) > 0 )
        write( f2, buf, n );
```

```
    exit( 0 );  
}
```

open is rather like **fopen** except that instead of a buffer, its second argument tells what sort of access is required. 0 implies reading, 1 implies writing, and 2 is both read and write. A -1 is returned if any error takes place, for example if the file doesn't exist.

creat is similar to **fcreat**, except that the second argument is the protection mode in which the output file is to be created. "0666" is read and write permission for everyone. **creat** opens the file for writing only.

I/O is done by **read** and **write**. In both cases, the first argument is the file descriptor returned by a previous **open** or **creat**. The second argument is the place where the data is to come from or go to. The third argument is the number of bytes to be transferred. On reading, the actual number of bytes returned may be less than this value. Zero bytes implies end of file, a -1 implies an error of some sort. For writing, the returned value is the number of bytes actually written; it is generally an error if this isn't equal to the number supposed to be written.

The routine **close** may be used to close files after you are done with them. Termination of the program via **exit** or return from the main program closes all files, but since only 15 files can be open at once, **close** has to be used when many files are read or written.

It is instructive to see how **read** and **write** can be used to construct special versions of things like **getchar**, **putchar**, etc. For example, if you don't want buffered input, you have to write your own **getchar**:

```
getchar() {  
    int c, n;  
    c = 0;  
    n = read( 0, &c, 1 );  
    return( n>0 ? c : 0 );  
}
```

At this point you might wonder why unbuffered input should be useful. In practice, it usually doesn't matter, but consider the following sequence of commands:

```
ed stuff  
1,$ s/ */ /g  
w stuff  
q  
opr stuff
```

which uses the editor to replace all sequences of multiple blanks in "stuff" by a single blank, then prints "stuff" off-line. Now imagine this pair of commands in a shell command file, and finally consider what would happen if the editor read its standard input with buffering. Since the command file contains commands to both the editor and the shell, the editor, if it read with a large enough buffer, would consume not only not only its own commands but also the following *opr* command which is intended for the shell. One can imagine complicated schemes for pushing back unwanted input, but the simplest approach for the editor, and analogous programs, is to read a single character at a time so they can't disturb input not intended for them. (Incidentally, the portable library's **ungetc** function doesn't solve this problem, since it just pushes input back into an internal buffer which vanishes along with its caller on termination.)

Similarly you can use the fact that **printf** calls **putchar** explicitly to make your own error-printing routines. If you provide a **putchar** like this

```
putchar(c) {
    write(2, &c, 1);
    return(c);
}
```

then calls to `printf` will write on the terminal, unbuffered.

As another example, this code duplicates the `putchar` and `flush` routines described above except that it always buffers:

```
char  buf[512];
int   fdes  1;
int   nleft 512;
char  *nextfree  &buf[0];

putchar( c )
{
    *nextfree++ = c;
    if ( --nleft <= 0 )
        flush( );
}

flush( )
{
    write( fdes, buf, nextfree - buf );
    nleft = 512;
    nextfree = buf;
}
```

6.2. Random Access

The `seek` routine provides a way to move around in a file without actually reading or writing.

```
seek( fd, offset, ptr );
```

forces the current position in the file whose descriptor is `fd` to move to position `offset`, which is taken with respect to the location in the file specified by `ptr`. `ptr` can be 0, 1, or 2 to specify an offset measured from the beginning, from the current position, or from the end of the file respectively. For example, to append to a newly-opened file,

```
seek( fd, 0, 2 );
```

and to get back to the beginning ("rewind"),

```
seek( fd, 0, 0 );
```

With `seek`, it is possible to treat files more or less like large arrays, at the price of slower access. Here is a routine to read an arbitrary record from an arbitrary place in a file.

```
get( fd, pos, buf, n )
int fd, pos, n;
char *buf;
{
    seek( fd, pos, 0 );    /* get to pos */
    n = read( fd, buf, n );
    return( n );
}
```

Since integers have only 16 bits, the **offset** specified is limited to 65,536; for this reason, **ptr** values of 3, 4, 5 cause **seek** to multiply the given offset by 512 (the number of bytes in one physical block) and then interpret **ptr** as if it were 0, 1, or 2 respectively. Thus to get to an arbitrary place in a large file you need two seeks, first one which selects the block, then one which has **ptr** equal to 1 and moves to the desired byte within the block.

6.3. Error Processing

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

7. Executing Commands From Programs

It is often easier to use a program written by someone else instead of inventing one's own. This section describes how to call a command from within a running program.

7.1. With the Portable C Library

The portable C library routine **system** takes one argument, a command string exactly as you would have typed it at the terminal (except for the new-line at the end) and executes it. This is probably the easiest way to execute a command from within a running program. For instance, to time-stamp the output of a program,

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

If the command string has to be built from pieces, the core-to-core formatting capabilities of **printf** may be useful (in the portable library only).

7.2. With Standard UNIX

If you're not using the portable C library, or if you need finer control over what happens, you will have to construct calls to other programs using the more primitive routines that the portable library's **system** routine is based on. For no good reason, the standard library doesn't have an equivalent of **system**.

First, you can execute another program *without returning*, by using the routine **execl**, (described under *exec* in Section II of the manual). To print the date as the last action of a running program, you can say

```
execl( "/bin/date", "date", 0 );
```

The first argument to **execl** is the *file name* of the command; you have to know where it's found. The second argument is conventionally the program name (that is, the last component of the file name), but this is seldom used except as a place-holder. If the command takes arguments, they are strung out after this, and the whole list is followed by a 0 to terminate it.

The `exec1` call overlays your program with `date`, runs it, then exits. More realistically, your program might fall into two or more phases that communicate only through temporary files, like the assembler. Here it is natural to make the second pass simply an `exec1` call from the first.

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

```
exec1( "/bin/date", "date", 0 );
exec1( "/usr/bin/date", "date", 0 );
printf( "Someone stole 'date'\n", 0 );
```

Another version of `exec`, `execv`, is useful when you don't know in advance how many arguments there are going to be. The call is

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

where `argp` is a list of pointers to the arguments; the last pointer must be followed by a 0 so `execv` can tell where the list ends. As with `exec1`, `filename` is the file in which the program is found, and `argp[0]` is the name of the program.

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

```
exec1( "/bin/sh", "sh", "-c", commandline, 0 );
```

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

7.3. Regaining Control

So far what we've talked about isn't really all that useful by itself. Now we show how to regain control after running a program with `exec1` or `execv`. Since these routines simply overlay the new program on the old one, to save the old one requires that it first be split into two copies; one of these can be overlaid, while the other waits for the new, overlaying program to finish. The splitting is done by the routine called `fork`:

```
pid = fork( );
```

splits the program into two copies, both of which continue to run. The only difference between the two is the value of `pid`. In one of these processes (the "child"), `pid` is zero; in the other (the "parent"), `pid` is non-zero; it is the process name of the child. Thus the basic way to call, and return from, another program is

```
if ( (fork( ) == 0 )
    exec1( "/bin/sh", "sh", "-c", command, 0 );
```

And in fact, except for handling errors, this is sufficient. The `fork` makes two copies of your program. In the child, the value returned by `fork` is zero, so it calls `exec1` which does the `command` and then dies. In the parent, `fork` returns non-zero so it skips the `exec1`.

More often, the parent wants to wait for the child to terminate, so output doesn't get scrambled. This can be done with

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

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

The status return word set by `wait` encodes in its low-order byte the system's idea of the child's termination status; it is 0 for OK and non-zero to indicate various kinds of problems like those mentioned in section 8. The high-order byte is taken from the argument of the call to `[c]exit` which caused a normal termination of the child process. At the moment, the standard command interpreter (the shell) isn't fussy about termination status, but it is good coding practice for all programs to return meaningful status; someday, they may be called by another program which cares whether they worked right.

When your program is called by the shell, the three file descriptors 0, 1, and 2 are set up pointing at the right files, and all other possible file descriptors are available for use. When you call another program, correct etiquette suggests making sure the same conditions hold. Neither of the `exec` calls affects open files in any way. Remember too, that both `fork` and `exec` create processes whose address space is distinct from that of their caller. Some buffer-flushing may be needed before using these calls. Conversely, if a caller buffers an input stream, the callee will lose the read-ahead information. (Essentially the same syndrome was discussed in §6.1.)

8. Signals — Interrupts and all that

This section is concerned with how you can make your program deal gracefully with signals from the outside world (like interrupts), and with program faults. Since there's nothing very useful that can be done from within C about program faults, which arise mainly from illegal memory references or from execution of peculiar instructions, we'll discuss only the outside-world signals: `interrupt`, which is sent when the DEL character is typed; `quit`, generated by the FS character; and `hangup`, caused by hanging up the phone. When one of these events occurs, the signal is sent to all processes which were started from the corresponding typewriter; unless other arrangements have been made, the signal terminates the process. In the `quit` case, a core image file is written, usually for debugging purposes.

The routine which alters the default action is `signal`, described in section II of [1]. It has two arguments: the first names the signal, the second specifies how to treat it. If the second argument is 1, the signal is ignored; if it is 0, the default action is restored. Thus

```
#define SIGINT 2
...
signal(SIGINT, 1);
```

ignores interrupts, while

```
signal(SIGINT, 0);
```

restores the default action of process termination. Such coding is seldom needed (though see below) because there is a command which runs another program with these three signals ignored:

```
nohup program &
```

runs `program` (with arguments if you like) in such a way that you can hang up on it without fear. Usually you would follow the command with an "&"; otherwise your terminal will be

firmly tied up. If the command is going to run for a long time, you might also use `nice`:

```
nice nohup program &
```

`Nice` lowers the priority of `program` so it won't hog the machine. Incidentally, starting a command with the "&" automatically causes interrupts and quits to be ignored, so you can compute in the background, and edit and debug in the foreground without danger. However, hangups will still terminate "&" programs.

Finally, the second argument to `signal` may be the name of a function (which, incidentally, has to be declared explicitly if the compiler hasn't seen it already). In this case, the named routine will be called when the signal occurs. Most commonly this facility is used to allow the program to clean up unfinished business before terminating, for example to delete a temporary file:

```
main( )
{
    extern int onintr( );
    if ( (signal(SIGINT, 1) & 1) == 0 )
        signal( SIGINT, onintr );
    /* Process ... */
}

onintr( )
{
    unlink( tempfile );
    exit( 100 );
}
```

Why the test and the double call to `signal`? It's quite simple: suppose this "interactive" program were run non-interactively, say under "&" with input from a file. If it began by announcing that all interrupts were to be sent to the `onintr` routine, that would in fact occur, even if the user meant to interrupt only some foreground process he happened to be running. The code as written depends on the fact that `signal` returns the previous value of its argument; the if clause asks whether interrupt was previously being ignored (value is odd, e.g. 1) and only if not does it request the call to `onintr`. In other words, if interrupts were being ignored when the program was called, they should still be ignored.

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

```
main( ) {
    extern int onintrup( );
    setexit( );
    if ( (signal(SIGINT, 1) & 1) == 0 )
        signal( SIGINT, onintrup );
    for ( ; ) {
        /* main processing loop */
    }
}
```

```
onintrup( )
{
    printf( "\nInterrupt\n" );
    reset( );
}
```

When an interrupt occurs, a call is forced to the `onintrup` routine, which can print a message (and perhaps set flags, etc.). `reset` is a non-local goto to the location after the last call to `setexit`, so control (and the stack level) will pop back to the place in the main routine where the signal is set up and the main loop entered. Notice, by the way, that `signal` gets called again after an interrupt occurs. This is necessary; most signals are automatically reset to their default action when they occur.

Some programs which want to detect signals simply can't be stopped at an arbitrary point, for example in the middle of updating a linked list. If the routine called on occurrence of a signal sets a flag and then returns instead of calling `exit` or `reset`, execution will continue at the exact point it was interrupted. The interrupt flag can then be tested at some convenient point in the main loop.

There is one difficulty associated with this approach. Suppose the program is reading the typewriter when the interrupt is sent. The specified routine is duly called; it sets its flag and returns. If it were really true, as we said above, that "execution resumes at the exact point it was interrupted," the program would continue reading the typewriter until the user typed another line. This behavior might well be confusing, since the user might not know that the program is reading; he presumably would prefer to have the signal take effect instantly. The method chosen to resolve this difficulty is to terminate the typewriter read when execution resumes after the signal, returning an error code which indicates what happened.

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

```
if ( getchar( ) == '\0' )
    if ( intflag )
        /* interrupt processing */
    else
        /* end-of-file or error processing */
```

A final subtlety to keep in mind becomes important when signal-catching is combined with execution of other programs. Suppose your program catches interrupts, and also includes a method (like "!" in the editor) whereby other programs can be executed. Your code should look something like this:

```
signal( SIGINT, 1 ); /* ignore interrupts */
if ( fork( ) == 0 )
    execl( ... );
wait( ... );
signal( SIGINT, onintrup ); /* restore interrupts */
```

Why is this? Again, it's not obvious but not really difficult. Suppose the program you call catches its own interrupts. If you interrupt the subprogram, it will get the signal and return to its main loop, and probably read your typewriter. But so also will the calling program pop out of its wait for the subprogram, and also read your typewriter. Having two processes reading

your typewriter is very unfortunate, since the system figuratively flips a coin to decide who should get each line of input. A simple way out is to have the parent program ignore interrupts until the child is done.

References

- [1] K. Thompson and D. M. Ritchie, *UNIX Programmer's Manual*, Sixth Edition. Bell Laboratories, 1975.
- [2] B. W. Kernighan, *Programming in C — A Tutorial*. In *The Programming Language C*, Bell Laboratories Computer Science Technical Report 31 (1975).
- [3] B. W. Kernighan, *UNIX For Beginners*. Bell Laboratories internal memorandum.
- [4] M. E. Lesk, *The Portable C Library*. In *The Programming Language C*, Bell Laboratories Computer Science Technical Report 31 (1975).







C Reference Manual

by Dennis M. Ritchie

Bell Laboratories
Murray Hill, New Jersey 07974

May 1, 1977

1. Introduction

C is a computer language which offers a rich selection of operators and data types and the ability to impose useful structure on both control flow and data. All the basic operations and data objects are close to those actually implemented by most real computers, so that a very efficient implementation is possible, but the design is not tied to any particular machine and with a little care it is possible to write easily portable programs.

This manual describes the current version of the C language as it exists on the PDP-11, the Honeywell 6000, the IBM System/370, and the Interdata 8/32. Where differences exist, it concentrates on the PDP-11, but tries to point out implementation-dependent details. With few exceptions, these dependencies follow directly from the underlying properties of the hardware; the various compilers are generally quite compatible.

2. Lexical conventions

Blanks, tabs, newlines, and comments as described below are ignored except as they serve to separate tokens. Some space is required to separate otherwise adjacent identifiers, keywords, and constants.

If the input stream has been parsed into tokens up to a given character, the next token is taken to include the longest string of characters which could possibly constitute a token.

2.1 Comments

The characters `/*` introduce a comment, which terminates with the characters `*/`. Comments do not nest.

2.2 Identifiers (Names)

An identifier is a sequence of letters and digits; the first character must be alphabetic. The underscore `'_'` counts as alphabetic. Upper and lower case letters are considered different. On the PDP-11, no more than the first eight characters are significant, and only the first seven for external identifiers.

2.3 Keywords

The following identifiers are reserved for use as keywords, and may not be used otherwise:

Warning: The data type name `short` is *not* recognized by the version of the C compiler that is distributed as part of PWB/UNIX Edition 1.0.

int	extern	else
char	register	for
float	typedef	do
double	static	while
struct	goto	switch
union	return	case
long	sizeof	default
short	break	entry
unsigned	continue	
auto	if	

The `entry` keyword is not currently implemented by any compiler but is reserved for future use. Some implementations also reserve the word `fortran`.

2.4 Constants

There are several kinds of constants, as follows:

2.4.1 Integer constants

An integer constant consisting of a sequence of digits is taken to be octal if it begins with 0 (digit zero), decimal otherwise. The digits 8 and 9 have octal value 10 and 11 respectively. A sequence of digits preceded by 0x or 0X (digit zero) is taken to be a hexadecimal integer. The hexadecimal digits include a or A through f or F with values 10 through 15. A decimal constant whose value exceeds the largest signed machine integer (32767 on the PDP-11) is taken to be long; an octal or hex constant which exceeds the largest unsigned machine integer (0177777 or 0xFFFF on the PDP-11) is likewise taken to be long.

2.4.2 Explicit long constants

A decimal, octal, or hexadecimal integer constant immediately followed by l (letter ell) or L is a long constant, which, on the PDP-11, has 32 significant bits. As discussed below, on other machines integer and long values may be considered identical.

2.4.3 Character constants

A character constant is a sequence of characters enclosed in single quotes `'`. Within a character constant a single quote must be preceded by a backslash `\`. Certain non-graphic characters, and `\` itself, may be escaped according to the following table:

BS	<code>\b</code>
NL (LF)	<code>\n</code>
CR	<code>\r</code>
HT	<code>\t</code>
FF	<code>\f</code>
<i>ddd</i>	<code>\ddd</code>
<code>\</code>	<code>\\</code>

The escape `\ddd` consists of the backslash followed by 1, 2, or 3 octal digits which are taken to specify the value of the desired character. A special case of this construction is `\0` (not followed by a digit) which indicates the character NUL. If the character following a backslash is not one of those specified, the backslash vanishes.

The value of a single-character constant is the numerical value of the character in the machine's character set (ASCII for the PDP-11). On the PDP-11 at most two characters are permitted in a character constant and the second character of a pair is stored in the high-order byte of the integer value. Character constants with more than one character are inherently machine-dependent and should be avoided.

2.4.4 Floating constants

A floating constant consists of an integer part, a decimal point, a fraction part, an *e* or *E*, and an optionally signed integer exponent. The integer and fraction parts both consist of a sequence of digits. Either the integer part or the fraction part (not both) may be missing; either the decimal point or the *e* and the exponent (not both) may be missing. Every floating constant is taken to be double-precision.

2.5 Strings

A string is a sequence of characters surrounded by double quotes `"`. A string has type `'array of characters'` and storage class `'static'` (see below) and is initialized with the given characters. The compiler places a null byte `'\0'` at the end of each string so that programs which scan the string can find its end. In a string, the character `"` must be preceded by a `'\'`; in addition, the same escapes as described for character constants may be used. Finally, a `'\'` and an immediately following new-line are ignored.

All strings, even when written identically, are distinct.

3. Syntax notation

In the syntax notation used in this manual, syntactic categories are indicated by *italic* type, and literal words and characters in *sans-serif* type. Alternatives are listed on separate lines. An optional terminal or non-terminal symbol is indicated by the subscript `'opt.'` so that

$$\{ expression_{opt} \}$$

would indicate an optional expression in braces. The complete syntax is given in §16, in the notation of YACC.

4. What's in a Name?

C bases the interpretation of an identifier upon two attributes of the identifier: its *storage class* and its *type*. The storage class determines the location and lifetime of the storage associated with an identifier; the type determines the meaning of the values found in the identifier's storage.

There are four declarable storage classes: automatic, static, external, and register. Automatic variables are local to each invocation of a block, and are discarded upon exit from the block; static variables are local to a block, but retain their values upon reentry to a block even after control has left the block; external variables exist and retain their values throughout the execution of the entire program, and may be used for communication between functions, even separately compiled functions. Register variables are (if possible) stored in the fast registers of the machine; like automatic variables they are local to each block and disappear on exit from the block.

C supports several fundamental types of objects:

Objects declared as characters (`char`) are large enough to store any member of the implementation's character set, and if a genuine character is stored in a character variable, its value is equivalent to the integer code for that character. Other quantities may be stored into character variables, but the implementation is machine-dependent. On the PDP-11, characters are stored as signed 8-bit integers, and the character set is ASCII.

Up to three sizes of integer, declared `short int`, `int`, and `long int` are available. Longer integers provide no less storage than shorter ones, but the implementation may make either short integers, or long integers, or both equivalent to plain integers. 'Plain' integers have the natural size suggested by the host machine architecture; the other sizes are provided to meet special needs. On the PDP-11, short and plain integers are both represented in 16-bit 2's complement notation. Long integers are 32-bit 2's complement.

Unsigned integers, declared `unsigned`, obey the laws of arithmetic modulo 2^n where n is the number of bits in the representation. (16 on the PDP-11; long and short unsigned quantities

are not supported.)

Single precision floating point (*float*) quantities, on the PDP-11, have magnitude in the range approximately $10^{\pm 38}$ or 0; their precision is 24 bits or about seven decimal digits.

Double-precision floating-point (*double*) quantities on the PDP-11 have the same range as floats and a precision of 56 bits or about 17 decimal digits. Some implementations may make float and double synonymous.

Because objects of these types can usefully be interpreted as numbers, they will be referred to as *arithmetic* types. Types *char* and *int* of all sizes will collectively be called *integral* types. *Float* and *double* will collectively be called *floating* types.

Besides the fundamental arithmetic types there is a conceptually infinite class of derived types constructed from the fundamental types in the following ways:

arrays of objects of most types;

functions which return objects of a given type;

pointers to objects of a given type;

structures containing a sequence of objects of various types;

unions capable of containing any one of several objects of various types.

In general these methods of constructing objects can be applied recursively.

5. Objects and lvalues

An *object* is a manipulatable region of storage; an *lvalue* is an expression referring to an object. An obvious example of an lvalue expression is an identifier. There are operators which yield lvalues: for example, if *E* is an expression of pointer type, then **E* is an lvalue expression referring to the object to which *E* points. The name 'lvalue' comes from the assignment expression '*E1 = E2*' in which the left operand *E1* must be an lvalue expression. The discussion of each operator below indicates whether it expects lvalue operands and whether it yields an lvalue.

6. Conversions

A number of operators may, depending on their operands, cause conversion of the value of an operand from one type to another. This section explains the result to be expected from such conversions. §6.6 summarizes the conversions demanded by most ordinary operators; it will be supplemented as required by the discussion of each operator.

6.1 Characters and integers

A character or a short integer may be used wherever an integer may be used. In all cases the value is converted to an integer. Conversion of a short integer always involves sign extension; short integers are signed quantities. Whether or not sign-extension occurs for characters is machine dependent, but it is guaranteed that a member of the standard character set is non-negative. On the PDP-11, character variables range in value from -128 to 127; a character constant specified using an octal escape also suffers sign extension and may appear negative, for example `'\214'`

When a longer integer is converted to a shorter or to a char, it is truncated on the left.

6.2 Float and double

All floating arithmetic in C is carried out in double-precision; whenever a float appears in an expression it is lengthened to double by zero-padding its fraction. When a double must be converted to float, for example by an assignment, the double is rounded before truncation to float length.

6.3 Floating and integral

Conversions of floating values to integral type tend to be rather machine-dependent. On the PDP-11, truncation is towards 0. The result is undefined if the value will not fit in the space provided.

Conversions of integral values to floating type are well behaved. Some loss of precision occurs if the destination lacks sufficient bits.

6.4 Pointers and integers

An integer or long integer may be added to or subtracted from a pointer; in such a case the first is converted as specified in the discussion of the addition operator.

Two pointers to objects of the same type may be subtracted; in this case the result is converted to an integer as specified in the discussion of the subtraction operator.

6.5 Unsigned

Whenever an unsigned integer and a plain integer are combined, the plain integer is converted to unsigned and the result is unsigned. The value (on the PDP-11) is the least unsigned integer congruent to the signed integer (modulo 2^{16}). Because of the 2's complement notation, this conversion is conceptual and there is no actual change in the bit pattern.

When an unsigned integer is converted to long, the value of the result is the same numerically as that of the unsigned integer. Thus the conversion amounts to padding with zeros on the left.

6.6 Arithmetic conversions

A great many operators cause conversions and yield result types in a similar way. This pattern will be called the 'usual arithmetic conversions.'

First, any operands of type `char` or `short` are converted to `int`, and any of type `float` are converted to `double`.

Then, if either operand is `double`, the other is converted to `double` and that is the type of the result.

Otherwise, if either operand is `long`, the other is converted to `long` and that is the type of the result.

Otherwise, if either operand is `unsigned`, the other is converted to `unsigned` and that is the type of the result.

Otherwise, both operands must be `int`, and that is the type of the result.

7. Expressions

The precedence of expression operators is the same as the order of the major subsections of this section (highest precedence first). Thus the expressions referred to as the operands of `+` (§7.4) are those expressions defined in §§7.1-7.3. Within each subsection, the operators have the same precedence. Left- or right-associativity is specified in each subsection for the operators discussed therein. The precedence and associativity of all the expression operators is summarized in the collected grammar.

Otherwise the order of evaluation of expressions is undefined. In particular the compiler considers itself free to compute subexpressions in the order it believes most efficient, even if the subexpressions involve side effects. Expressions involving a commutative and associative operator may be rearranged arbitrarily, even in the presence of parentheses; to force a particular order of evaluation an explicit temporary must be used.

7.1 Primary expressions

Primary expressions involving `.`, `>`, subscripting, and function calls group left to right.

```
primary-expression:
    identifier
    constant
    string
    ( expression )
    primary-expression [ expression ]
    primary-expression ( expression-listopt )
    primary-lvalue . identifier
    primary-expression -> identifier
```

```
expression-list:
    expression
    expression-list , expression
```

An identifier is a primary expression, provided it has been suitably declared as discussed below. Its type is specified by its declaration. However, if the type of the identifier is 'array of ...', then the value of the identifier-expression is a pointer to the first object in the array, and the type of the expression is 'pointer to ...'. Moreover, an array identifier is not an lvalue expression. Likewise, an identifier which is declared 'function returning ...', when used except in the function-name position of a call, is converted to 'pointer to function returning ...'.

A decimal, octal, character, or floating constant is a primary expression. Its type may be int, long, or double depending on its form.

A string is a primary expression. Its type is originally 'array of char'; but following the same rule given above for identifiers, this is modified to 'pointer to char' and the result is a pointer to the first character in the string. (There is an exception in certain initializers; see §8.6.)

A parenthesized expression is a primary expression whose type and value are identical to those of the unadorned expression. The presence of parentheses does not affect whether the expression is an lvalue.

A primary expression followed by an expression in square brackets is a primary expression. The intuitive meaning is that of a subscript. Usually, the primary expression has type 'pointer to ...', the subscript expression is int, and the type of the result is '...'. The expression 'E1[E2]' is identical (by definition) to '*((E1) + (E2))'. All the clues needed to understand this notation are contained in this section together with the discussions in §§ 7.1, 7.2, and 7.4 on identifiers, *, and + respectively; §14.3 below summarizes the implications.

A function call is a primary expression followed by parentheses containing a possibly empty, comma-separated list of expressions which constitute the actual arguments to the function. The primary expression must be of type 'function returning ...', and the result of the function call is of type '...'. As indicated below, a hitherto unseen identifier followed immediately by a left parenthesis is contextually declared to represent a function returning an integer; thus in the most common case, integer-valued functions need not be declared.

Any actual arguments of type float are converted to double before the call; any of type char or short are converted to int.

In preparing for the call to a function, a copy is made of each actual parameter; thus, all argument-passing in C is strictly by value. A function may change the values of its formal parameters, but these changes cannot affect the values of the actual parameters. On the other hand, it is possible to pass a pointer on the understanding that the function may change the value of the object to which the pointer points. The order of evaluation of arguments is undefined by the language; take note that the various compilers differ.

Recursive calls to any function are permitted.

A primary expression followed by a dot followed by an identifier is an expression. The first expression must be an lvalue naming a structure or union, and the identifier must name a member of the structure or union. The result is an lvalue referring to the named member of the structure or union.

A primary expression followed by an arrow (built from a '-' and a '>') followed by an identifier is an expression. The first expression must be a pointer to a structure or a union and the identifier must name a member of that structure or union. The result is an lvalue referring to the named member of the structure or union to which the pointer expression points.

Thus the expression 'E1->MOS' is the same as '(*E1).MOS'. Structures and unions are discussed in §8.5. The rules given here for the use of structures and unions are not enforced strictly, in order to allow an escape from the typing mechanism. See §14.1.

7.2 Unary operators

Expressions with unary operators group right-to-left.

unary-expression:

- * *expression*
- & *lvalue*
- *expression*
- ! *expression*
- ~ *expression*
- ++ *lvalue*
- *lvalue*
- lvalue* ++
- lvalue* --
- (*type-name*) *expression*
- sizeof *expression*
- sizeof (*type-name*)

The unary * operator means *indirection*: the expression must be a pointer, and the result is an lvalue referring to the object to which the expression points. If the type of the expression is 'pointer to ...', the type of the result is '...'.
The result of the unary & operator is a pointer to the object referred to by the lvalue. If the type of the lvalue is '...', the type of the result is 'pointer to ...'.
The result of the unary - operator is the negative of its operand. The usual arithmetic conversions are performed. The negative of an unsigned quantity is computed by subtracting its value from 2ⁿ, where n is 16 on the PDP-11.
The result of the logical negation operator ! is 1 if the value of its operand is 0, 0 if the value of its operand is non-zero. The type of the result is int. It is applicable to any arithmetic type or to pointers.
The ~ operator yields the one's complement of its operand. The usual arithmetic conversions are performed. The type of the operand must be integral.
The object referred to by the lvalue operand of prefix '++' is incremented. The value is the new value of the operand, but is not an lvalue. The expression '++a' is equivalent to '(a += 1)'. See the discussions of addition (§7.4) and assignment operators (§7.14) for information on conversions.
The lvalue operand of prefix '--' is decremented analogously to the ++ operator.
When postfix '++' is applied to an lvalue the result is the value of the object referred to by the lvalue. After the result is noted, the object is incremented in the same manner as for the prefix ++ operator. The type of the result is the same as the type of the lvalue expression.
When postfix '--' is applied to an lvalue the result is the value of the object referred to by the lvalue. After the result is noted, the object is decremented in the manner as for the

prefix `---` operator. The type of the result is the same as the type of the lvalue expression.

An expression preceded by the parenthesized name of a data type causes conversion of the value of the expression to the named type. The construction of type names is described in §8.7.

The `sizeof` operator yields the size, in bytes, of its operand. (A *byte* is undefined by the language except in terms of the value of `sizeof`. However in all existing implementations a byte is the space required to hold a `char`.) When applied to an array, the result is the total number of bytes in the array. The size is determined from the declarations of the objects in the expression. This expression is semantically an integer constant and may be used anywhere a constant is required. Its major use is in communication with routines like storage allocators and I/O systems.

The `sizeof` operator may also be applied to a parenthesized type name. In that case it yields the size, in bytes, of an object of the indicated type.

The construction `'sizeof(type)'` is taken to be a unit, so the expression `'sizeof(type)-2'` is the same as `'(sizeof(type))-2'`.

7.3 Multiplicative operators

The multiplicative operators `*`, `/`, and `%` group left-to-right. The usual arithmetic conversions are performed.

multiplicative-expression:
*expression * expression*
expression / expression
expression % expression

The binary `*` operator indicates multiplication. The `*` operator is associative and expressions with several multiplications at the same level may be rearranged.

The binary `/` operator indicates division. When positive integers are divided truncation is toward 0, but the form of truncation is machine-dependent if either operand is negative. In all cases it is true that $(a/b)*b + a\%b = a$. On the PDP-11, the remainder has the same sign as the dividend.

The binary `%` operator yields the remainder from the division of the first expression by the second. The usual arithmetic conversions are performed. On the PDP-11, the remainder has the same sign as the dividend. The operands must not be floating.

7.4 Additive operators

The additive operators `+` and `-` group left-to-right. The usual arithmetic conversions are performed. There are some additional type possibilities for each operator.

additive-expression:
expression + expression
expression - expression

The result of the `'+'` operator is the sum of the operands. A pointer to an object in an array and a value of any integral type may be added. The latter is in all cases converted to an address offset by multiplying it by the length of the object to which the pointer points. The result is a pointer of the same type as the original pointer, and which points to another object in the same array, appropriately offset from the original object. Thus if `P` is a pointer to an object in an array, the expression `'P+1'` is a pointer to the next object in the array.

No further type combinations are allowed.

The `+` operator is associative and expressions with several additions at the same level may be rearranged.

The result of the `'-'` operator is the difference of the operands. The usual arithmetic conversions are performed. Additionally, a value of any integral type may be subtracted from a

pointer, and then the same conversions as for addition apply.

If two pointers to objects of the same type are subtracted, the result is converted (by division by the length of the object) to an int representing the number of objects separating the pointed-to objects. This conversion will in general give unexpected results unless the pointers point to objects in the same array, since pointers, even to objects of the same type, do not necessarily differ by a multiple of the object-length.

7.5 Shift operators

The shift operators `<<` and `>>` group left-to-right. Both perform the usual arithmetic conversions on their operands, each of which must be integral. Then the right operand is converted to int; the type of the result is that of the left operand. The result is undefined if the right operand is negative or larger than the number of bits in the object.

shift-expression:

expression << expression
expression >> expression

The value of 'E1<<E2' is E1 (interpreted as a bit pattern) left-shifted E2 bits; vacated bits are 0-filled. The value of 'E1>>E2' is E1 right-shifted E2 bit positions. The shift is guaranteed to be logical (0-fill) if E1 is unsigned; otherwise it may be (and is, on the PDP-11) arithmetic (fill by a copy of the sign bit).

7.6 Relational operators

The relational operators group left-to-right, but this fact is not very useful; 'a<b<c' does not mean what it seems to.

relational-expression:

expression < expression
expression > expression
expression <= expression
expression >= expression

The operators `<` (less than), `>` (greater than), `<=` (less than or equal to) and `>=` (greater than or equal to) all yield 0 if the specified relation is false and 1 if it is true. The type of the result is int. The usual arithmetic conversions are performed. Two pointers may be compared, and the result depends on the relative locations in the address space of the pointed-to objects. Pointer comparison is portable only when the pointers point to objects in the same array.

7.7 Equality operators

equality-expression:

expression == expression
expression != expression

The `==` (equal to) and the `!=` (not equal to) operators are exactly analogous to the relational operators except for their lower precedence. (Thus 'a<b == c<d' is 1 whenever a<b and c<d have the same truth-value).

A pointer may be compared to an integer, but the result is machine dependent unless the integer is the constant 0. A pointer to which 0 has been assigned is guaranteed not to point to any object, and will appear to be equal to 0; in conventional usage, such a pointer is considered to be null.

7.8 Bitwise *and* operator

and-expression:

expression & expression

The **&** operator is associative and expressions involving **&** may be rearranged. The usual arithmetic conversions are performed; the result is the bit-wise 'and' function of the operands. The operator applies only to integral operands.

7.9 Bitwise *exclusive or* operator

exclusive-or-expression:
expression ^ expression

The **^** operator is associative and expressions involving **^** may be rearranged. The usual arithmetic conversions are performed; the result is the bit-wise 'exclusive or' function of the operands. The operator applies only to integral operands.

7.10 Bitwise *inclusive or* operator

inclusive-or-expression:
expression | expression

The **|** operator is associative and expressions with **|** may be rearranged. The usual arithmetic conversions are performed; the result is the bit-wise 'inclusive or' function of its operands. The operator applies only to integral operands.

7.11 Logical *and* operator

logical-and-expression:
expression && expression

The **&&** operator groups left-to-right. It returns 1 if both its operands are non-zero, 0 otherwise. Unlike **&**, **&&** guarantees left-to-right evaluation; moreover the second operand is not evaluated if the first operand is 0.

The operands need not have the same type, but each must have one of the fundamental types or be a pointer. The result is always int.

7.12 Logical *or* operator

logical-or-expression:
expression || expression

The **||** operator groups left-to-right. It returns 1 if either of its operands is non-zero, and 0 otherwise. Unlike **|**, **||** guarantees left-to-right evaluation; moreover, the second operand is not evaluated if the value of the first operand is non-zero.

The operands need not have the same type, but each must have one of the fundamental types or be a pointer. The result is always int.

7.13 Conditional operator

conditional-expression:
expression ? expression : expression

Conditional expressions group right-to-left. The first expression is evaluated and if it is non-zero, the result is the value of the second expression, otherwise that of third expression. If possible, the usual arithmetic conversions are performed to bring the second and third expressions to a common type; otherwise, if both are pointers of the same type, the result has the common type; otherwise, one must be a pointer and the other the constant 0, and the result has the type of the pointer. Only one of the second and third expressions is evaluated.

7.14 Assignment operators

There are a number of assignment operators, all of which group right-to-left. All require an lvalue as their left operand, and the type of an assignment expression is that of its left operand. The value is the value stored in the left operand after the assignment has taken place. The two parts of a compound assignment operator are separate tokens.

assignment-expression:

lvalue = *expression*
lvalue + = *expression*
lvalue - = *expression*
lvalue * = *expression*
lvalue / = *expression*
lvalue % = *expression*
lvalue >> = *expression*
lvalue << = *expression*
lvalue & = *expression*
lvalue ^ = *expression*
lvalue | = *expression*

Notice that the representation of the compound assignment operators has changed; formerly the '=' came first and the other operator came second (without any space). The compiler continues to accept the previous notation.

In the simple assignment with '=', the value of the expression replaces that of the object referred to by the lvalue. If both operands have arithmetic type, the right operand is converted to the type of the left preparatory to the assignment.

The behavior of an expression of the form 'E1 op = E2' may be inferred by taking it as equivalent to 'E1 = E1 op (E2)'; however, E1 is evaluated only once. In += and -=, the left operand may be a pointer, in which case the (integral) right operand is converted as explained in §7.4; all right operands and all non-pointer left operands must have arithmetic type.

The compiler currently allows a pointer to be assigned to an integer, an integer to a pointer, and a pointer to a pointer of another type. The assignment is a pure copy operation, with no conversion. This usage is nonportable, and may produce pointers which cause addressing exceptions when used. However, it is guaranteed that assignment of the constant 0 to a pointer will produce a null pointer distinguishable from a pointer to any object.

7.15 Comma operator

comma-expression:

expression , *expression*

A pair of expressions separated by a comma is evaluated left-to-right and the value of the left expression is discarded. The type and value of the result are the type and value of the right operand. This operator groups left-to-right. In contexts where comma is given a special meaning, for example in a list of actual arguments to functions (§7.1) and lists of initializers (§8.6), the comma operator as described in this section can only appear in parentheses; for example, 'f(a, (t = 3, t+2), c)' has three arguments, the second of which has the value 5.

8. Declarations

Declarations are used within function definitions to specify the interpretation which C gives to each identifier; they do not necessarily reserve storage associated with the identifier. Declarations have the form

declaration:

*decl-specifiers declarator-list*_{opt} ;

The declarators in the declarator-list contain the identifiers being declared. The decl-specifiers consist of a sequence of type and storage class specifiers.

decl-specifiers:

type-specifier decl-specifiers_{opt}
sc-specifier decl-specifiers_{opt}

The list must be self-consistent in a way described below.

8.1 Storage class specifiers

The sc-specifiers are:

sc-specifier:

auto
static
extern
register
typedef

The **typedef** specifier does not reserve storage and is called a 'storage class specifier' only for syntactic convenience; it is discussed in §8.8.

The meanings of the various storage classes were discussed in §4.

The **auto**, **static**, and **register** declarations also serve as definitions in that they cause an appropriate amount of storage to be reserved. In the **extern** case there must be an external definition (§10) for the given identifiers somewhere outside the function in which they are declared.

A **register** declaration is best thought of as an **auto** declaration, together with a hint to the compiler that the variables declared will be heavily used. Only the first few (three, for the PDP-11) such declarations are effective. Moreover, only variables of certain types will be stored in registers; on the PDP-11, they are **int**, **char**, or **pointer**. One restriction applies to **register** variables: the address-of operator **&** cannot be applied to them. Smaller, faster programs can be expected if register declarations are used appropriately, but future developments may render them unnecessary.

At most one sc-specifier may be given in a declaration. If the sc-specifier is missing from a declaration, it is taken to be **auto** inside a function, **extern** outside. Exception: functions are always **extern**.

8.2 Type specifiers

The type-specifiers are

type-specifier:

char
short
int
long
unsigned
float
double
struct-or-union-specifier
typedef-name

The words **long**, **short**, and **unsigned** may be thought of as adjectives; the following combinations are acceptable (in any order).

short int
long int
unsigned int
long float

The meaning of the last is the same as double. Otherwise, at most one type-specifier may be given in a declaration. If the type-specifier is missing from a declaration, it is taken to be int.

Specifiers for structures and unions are discussed in §8.5; declarations with typedef names are discussed in §8.8.

8.3 Declarators

The declarator-list appearing in a declaration is a comma-separated sequence of declarators, each of which may have an initializer.

declarator-list:
init-declarator
init-declarator , *declarator-list*

init-declarator:
declarator *initializer*_{opt}

Initializers are discussed in §8.6. The specifiers in the declaration indicate the type and storage class of the objects to which the declarators refer. Declarators have the syntax:

declarator:
identifier
(*declarator*)
* *declarator*
declarator ()
declarator [*constant-expression*_{opt}]

The grouping is the same as in expressions.

8.4 Meaning of declarators

Each declarator is taken to be an assertion that when a construction of the same form as the declarator appears in an expression, it yields an object of the indicated type and storage class. Each declarator contains exactly one identifier; it is this identifier that is declared.

If an unadorned identifier appears as a declarator, then it has the type indicated by the specifier heading the declaration.

A declarator in parentheses is identical to the unadorned declarator, but the binding of complex declarators may be altered by parentheses. See the examples below.

If a declarator has the form

* D

for D a declarator, then the contained identifier has the type 'pointer to ...', where '...' is the type which the identifier would have had if the declarator had been simply D.

If a declarator has the form

D ()

then the contained identifier has the type 'function returning ...', where '...' is the type which the identifier would have had if the declarator had been simply D.

A declarator may have the form

D [constant-expression]

or

D[]

Such declarators make the contained identifier have type 'array.' If the unadorned declarator *D* would specify a non-array of type '...', then the declarator 'D[i]' yields a 1-dimensional array with rank *i* of objects of type '...'. If the unadorned declarator *D* would specify an *n*-dimensional array with rank $i_1 \times i_2 \times \dots \times i_n$, then the declarator *D*[i_{n+1}] yields an (*n*+1)-dimensional array with rank $i_1 \times i_2 \times \dots \times i_n \times i_{n+1}$.

In the first case the constant expression is an expression whose value is determinable at compile time, and whose type is int. (Constant expressions are defined precisely in §15.) The constant expression of an array declarator may be missing only for the first dimension. This notation is useful when the array is external and the actual declaration, which allocates storage, is given elsewhere. The constant-expression may also be omitted when the declarator is followed by initialization. In this case the size is calculated from the number of initial elements supplied.

An array may be constructed from one of the basic types, from a pointer, from a structure or union, or from another array (to generate a multi-dimensional array).

Not all the possibilities allowed by the syntax above are actually permitted. The restrictions are as follows: functions may not return arrays, structures or functions, although they may return pointers to such things; there are no arrays of functions, although there may be arrays of pointers to functions. Likewise a structure may not contain a function, but it may contain a pointer to a function.

As an example, the declaration

```
int i, *ip, f(), *fip(), (*pfi)();
```

declares an integer *i*, a pointer *ip* to an integer, a function *f* returning an integer, a function *fip* returning a pointer to an integer, and a pointer *pfi* to a function which returns an integer. It is especially useful to compare the last two. The binding of '*fip()' is '(fip())', so that the declaration suggests, and the same construction in an expression requires, the calling of a function *fip*, and then using indirection through the (pointer) result to yield an integer. In the declarator '(*pfi)()', the extra parentheses are necessary, as they are also in an expression, to indicate that indirection through a pointer to a function yields a function, which is then called.

As another example,

```
float fa[17], *afp[17];
```

declares an array of float numbers and an array of pointers to float numbers. Finally,

```
static int x3d[3][5][7];
```

declares a static three-dimensional array of integers, with rank $3 \times 5 \times 7$. In complete detail, *x3d* is an array of three items: each item is an array of five arrays; each of the latter arrays is an array of seven integers. Any of the expressions 'x3d', 'x3d[i]', 'x3d[i][j]', 'x3d[i][j][k]' may reasonably appear in an expression. The first three have type 'array', the last has type int.

8.5 Structure and union declarations

A structure is an object consisting of a sequence of named members. Each member may have any type. A union is an object which may, at a given time, contain any one of several members. Structure and union specifiers have the same form.

structure-or-union-specifier:

```
struct-or-union { struct-decl-list }  
struct-or-union identifier { struct-decl-list }  
struct-or-union identifier
```

struct-or-union:
struct
union

The struct-decl-list is a sequence of declarations for the members of the structure or union:

struct-decl-list:
struct-declaration
struct-declaration struct-decl-list

struct-declaration:
type-specifier struct-declarator-list ;

struct-declarator-list:
struct-declarator
struct-declarator , struct-declarator-list

In the usual case, a struct-declarator is just a declarator for a member of a structure or union. A structure member may also consist of a specified number of bits. Such a member is also called a *field*; its length is set off from the field name by a colon.

struct-declarator:
declarator
declarator : constant-expression
: constant-expression

Within a structure, the objects declared have addresses which increase as their declarations are read left-to-right. Each non-field member of a structure begins on an addressing boundary appropriate to its type. On the PDP-11 the only requirement is that non-characters begin on a word boundary; therefore, there may be 1-byte, unnamed holes in a structure. Field members are packed into machine integers; they do not straddle words. A field which does not fit into the space remaining in a word is put into the next word. No field may be wider than a word. On the PDP-11, fields are assigned right-to-left.

A struct-declarator with no declarator, only a colon and a width, indicates an unnamed field useful for padding to conform to externally-imposed layouts. As a special case, an unnamed field with a width of 0 specifies alignment of the next field at a word boundary. The 'next field' presumably is a field, not an ordinary structure member, because in the latter case the alignment would have been automatic.

The language does not restrict the types of things that are declared as fields, but implementations are not required to support any but integer fields. Moreover, even int fields may be considered to be unsigned. On the PDP-11, fields are not signed and have only integer values.

A union may be thought of as a structure all of whose members begin at offset 0 and whose size is sufficient to contain any of its members. At most one of the members can be stored in a union at any time.

A structure or union specifier of the second form, that is, one of

```
struct identifier { struct-decl-list }  
union identifier { struct-decl-list }
```

declares the identifier to be the *structure tag* (or union tag) of the structure specified by the list. A subsequent declaration may then use the third form of specifier, one of

```
struct identifier  
union identifier
```

Structure tags allow definition of self-referential structures; they also permit the long part of the declaration to be given once and used several times. It is however absurd to declare a structure or union which contains an instance of itself, as distinct from a pointer to an instance of itself.

The names of members and tags may be the same as ordinary variables. However, names of tags and members must be mutually distinct.

Two structures may share a common initial sequence of members; that is, the same member may appear in two different structures if it has the same type in both and if all previous members are the same in both. (Actually, the compiler checks only that a name in two different structures has the same type and offset in both, but if preceding members differ the construction is nonportable.)

A simple example of a structure declaration is

```
struct tnode {
    char tword[20];
    int count;
    struct tnode *left;
    struct tnode *right;
};
```

which contains an array of 20 characters, an integer, and two pointers to similar structures. Once this declaration has been given, the following declaration makes sense:

```
struct tnode s, *sp;
```

which declares *s* to be a structure of the given sort and *sp* to be a pointer to a structure of the given sort. With these declarations, the expression

```
sp->count
```

refers to the *count* field of the structure to which *sp* points;

```
s.left
```

refers to the left subtree pointer of the structure *s*. Finally,

```
s.right->tword[0]
```

refers to the first character of the *tword* member of the right subtree of *s*.

8.6 Initialization

A declarator may specify an initial value for the identifier being declared. The initializer is preceded by '=', and consists of an expression or a list of values nested in braces.

initializer:

```
= expression
= { initializer-list }
= { initializer-list , }
```

initializer-list:

```
expression
initializer-list , initializer-list
{ initializer-list }
```

The '=' is a new addition to the syntax, intended to alleviate potential ambiguities. The current compiler allows it to be omitted when the rest of the initializer is a very simple expression (just a name, string, or constant) or when the rest of the initializer is enclosed in braces.

All the expressions in an initializer for a static or external variable must be constant expressions, which are described in §15, or expressions which reduce to the address of a previously declared variable, possibly offset by a constant expression. Automatic or register variables may be initialized by arbitrary expressions involving previously declared variables.

When an initializer applies to a *scalar* (a pointer or an object of arithmetic type), it consists of a single expression, perhaps in braces. The initial value of the object is taken from the expression; the same conversions as for assignment are performed.

When the declared variable is an *aggregate* (a structure or array) then the initializer consists of a brace-enclosed, comma-separated list of initializers for the members of the aggregate,

written in increasing subscript or member order. If the aggregate contains subaggregates, this rule applies recursively to the members of the aggregate. If there are fewer initializers in the list than there are members of the aggregate, then the aggregate is padded with 0's. It is not permitted to initialize unions or automatic aggregates. Currently, the PDP-11 compiler also forbids initializing fields in structures.

Braces may be elided as follows. If the initializer begins with a left brace, then the succeeding comma-separated list of initializers initialize the members of the aggregate; it is erroneous for there to be more initializers than members. If, however, the initializer does not begin with a left brace, then only enough elements from the list are taken to account for the members of the aggregate; any remaining members are left to initialize the next member of the aggregate of which the current aggregate is a part.

A final abbreviation allows a char array to be initialized by a string. In this case successive members of the string initialize the members of the array.

For example,

```
int x[] = { 1, 3, 5 };
```

declares and initializes *x* as a 1-dimensional array which has three members, since no size was specified and there are three initializers.

```
float y[4][3] = {  
    { 1, 3, 5 },  
    { 2, 4, 6 },  
    { 3, 5, 7 },  
};
```

is a completely-bracketed initialization: 1, 3, and 5 initialize the first row of the array *y*[0], namely *y*[0][0], *y*[0][1], and *y*[0][2]. Likewise the next two lines initialize *y*[1] and *y*[2]. The initializer ends early and therefore *y*[3] is initialized with 0. Precisely the same effect could have been achieved by

```
float y[4][3] = {  
    1, 3, 5, 2, 4, 6, 3, 5, 7,  
};
```

The initializer for *y* begins with a left brace, but that for *y*[0] does not, therefore 3 elements from the list are used. Likewise the next three are taken successively for *y*[1] and *y*[2]. Also,

```
float y[4][3] = {  
    { 1 }, { 2 }, { 3 }, { 4 }  
};
```

initializes the first column of *y* (regarded as a two-dimensional array) and leaves the rest 0.

Finally,

```
char msg[] = "Syntax error on line %s\n";
```

shows a character array whose members are initialized with a string.

8.7 Type names

In two contexts (to specify type conversions explicitly, and as an argument of `sizeof`) it is desired to supply the name of a data type. This is accomplished using a 'type name,' which in essence is a declaration for an object of that type which omits the name of the object.

```
type-name:  
    type-specifier abstract-declarator
```

abstract-declarator:

```
empty  
( abstract-declarator )  
* abstract-declarator  
abstract-declarator ( )  
abstract-declarator [ constant-expressionopt ]
```

To avoid ambiguity, in the construction

```
( abstract-declarator )
```

the abstract-declarator is required to be nonempty. Under this restriction, it is possible to identify uniquely the location in the abstract-declarator where the identifier would appear if the construction were a declarator in a declaration. The named type is then the same as the type of the hypothetical identifier. For example,

```
int  
int *  
int *[3]  
int (*)[3]
```

name respectively the types 'integer,' 'pointer to integer,' 'array of 3 pointers to integers,' and 'pointer to an array of 3 integers.' As another example,

```
int i;  
...  
sin( (double) i);
```

calls the *sin* routine (which accepts a **double** argument) with an argument appropriately converted.

8.8 Typedef

Declarations whose 'storage class' is **typedef** do not define storage, but instead define identifiers which can be used later as if they were type keywords naming fundamental or derived types. Within the scope of a declaration involving **typedef**, each of the identifiers appearing as part of any declarators therein become syntactically equivalent to type keywords naming the type associated with the identifiers in the way described in §8.4.

typedef-name:
identifier

For example, after

```
typedef int MILES, *KLICKSP;  
typedef struct { double re, im;} complex;
```

the constructions

```
MILES distance;  
extern KLICKSP metricp;  
complex z, *zp;
```

are all legal declarations; the type of *distance* is 'int', that of *metricp* is 'pointer to int,' and that of *z* is the specified structure. *Zp* is a pointer to such a structure.

Typedef does not introduce brand new types, only synonyms for types which could be specified in another way. Thus in the example above *distance* is considered to have exactly the same type as any other int variable.

9. Statements

Except as indicated, statements are executed in sequence.

9.1 Expression statement

Most statements are expression statements, which have the form

expression ;

Usually expression statements are assignments or function calls.

9.2 Compound statement, or block

So that several statements can be used where one is expected, the compound statement (also, and equivalently, called 'block') is provided:

compound-statement:
{ *declaration-list*_{opt} *statement-list*_{opt} }

declaration-list:
declaration
declaration declaration-list

statement-list:
statement
statement statement-list

If any of the identifiers in the declaration-list were previously declared, the outer declaration is pushed down for the duration of the block, at which time it resumes its force.

Any initializations of **auto** or **register** variables are performed each time the block is entered at the top. It is currently possible (but a bad practice) to transfer into a block; in that case the initializations are not performed. Initializations of **static** variables are performed only once when the program begins execution. Inside a block, external declarations do not reserve storage so initialization is not permitted.

9.3 Conditional statement

The two forms of the conditional statement are

if (*expression*) *statement*
if (*expression*) *statement* else *statement*

In both cases the expression is evaluated and if it is non-zero, the first substatement is executed. In the second case the second substatement is executed if the expression is 0. As usual the 'else' ambiguity is resolved by connecting an **else** with the last encountered **elseless if**.

9.4 While statement

The **while** statement has the form

while (*expression*) *statement*

The substatement is executed repeatedly so long as the value of the expression remains non-zero. The test takes place before each execution of the statement.

9.5 Do statement

The **do** statement has the form

do *statement* **while** (*expression*) ;

The substatement is executed repeatedly until the value of the expression becomes zero. The test takes place after each execution of the statement.

9.6 For statement

The for statement has the form

```
for ( expression-1opt ; expression-2opt ; expression-3opt ) statement
```

This statement is equivalent to

```
expression-1;  
while ( expression-2 ) {  
    statement  
    expression-3;  
}
```

Thus the first expression specifies initialization for the loop; the second specifies a test, made before each iteration, such that the loop is exited when the expression becomes 0; the third expression typically specifies an incrementation which is performed after each iteration.

Any or all of the expressions may be dropped. A missing expression-2 makes the implied while clause equivalent to 'while(1)'; other missing expressions are simply dropped from the expansion above.

9.7 Switch statement

The switch statement causes control to be transferred to one of several statements depending on the value of an expression. It has the form

```
switch ( expression ) statement
```

The usual arithmetic conversion is performed on the expression, but the result must be int. The statement is typically compound. Any statement within the statement may be labelled with one or more case prefixes as follows:

```
case constant-expression :
```

where the constant expression must be int. No two of the case constants in the same switch may have the same value. Constant expressions are precisely defined in §15.

There may also be at most one statement prefix of the form

```
default :
```

When the switch statement is executed, its expression is evaluated and compared with each case constant. If one of the case constants is equal to the value of the expression, control is passed to the statement following the matched case prefix. If no case constant matches the expression, and if there is a default prefix, control passes to the prefixed statement. If no case matches and if there is no default then none of the statements in the switch is executed.

Case and default prefixes in themselves do not alter the flow of control, which continues unimpeded across such prefixes. To exit from a switch, see break, §9.8.

Usually the statement that is the subject of a switch is compound. Declarations may appear at the head of this statement, initializations of automatic or register variables are ineffective.

9.8 Break statement

The statement

```
break ;
```

causes termination of the smallest enclosing while, do, for, or switch statement; control passes to the statement following the terminated statement.

9.9 Continue statement

The statement

```
continue ;
```

causes control to pass to the loop-continuation portion of the smallest enclosing **while**, **do**, or **for** statement; that is to the end of the loop. More precisely, in each of the statements

```
while ( ... ) {           do {           for ( ... ) {
    ...                   ...             ...
    contin:;              } while ( ... );  contin:;
}                          }                          }
```

a **continue** is equivalent to 'goto contin'. (Following the 'contin:' is a null statement, §9.13.)

9.10 Return statement

A function returns to its caller by means of the **return** statement, which has one of the forms

```
return ;
return expression ;
```

In the first case the returned value is undefined. In the second case, the value of the *expression* is returned to the caller of the function. If required, the *expression* is converted, as if by assignment, to the type of the function in which it appears. Flowing off the end of a function is equivalent to a return with no returned value.

9.11 Goto statement

Control may be transferred unconditionally by means of the statement

```
goto identifier ;
```

The *identifier* must be a label (§9.12) located in the current function. Previous versions of C had an incompletely implemented notion of label variable, which has been withdrawn.

9.12 Labelled statement

Any statement may be preceded by label prefixes of the form

```
identifier :
```

which serve to declare the *identifier* as a label. The only use of a label is as a target of a **goto**. The scope of a label is the current function, excluding any sub-blocks in which the same *identifier* has been redeclared. See §11.

9.13 Null statement

The null statement has the form

```
;
```

A null statement is useful to carry a label just before the '}' of a compound statement or to supply a null body to a looping statement such as **while**.

10. External definitions

A C program consists of a sequence of external definitions. An external definition declares an *identifier* to have storage class **extern** (by default) or perhaps **static**, and a specified type. The type-specifier (§8.2) may also be empty, in which case the type is taken to be **int**. The scope of external definitions persists to the end of the file in which they are declared just as the effect of declarations persists to the end of a block. The syntax of external definitions is the same as that of all declarations, except that only at this level may the code for functions be given.

10.1 External function definitions

Function definitions have the form

function-definition:
*decl-specifiers*_{opt} *function-declarator* *function-body*

The only sc-specifiers allowed among the decl-specifiers are **extern** or **static**; See §11.2 for the distinction between them. A function declarator is similar to a declarator for a 'function returning ...' except that it lists the formal parameters of the function being defined.

function-declarator:
declarator (*parameter-list*_{opt})

parameter-list:
identifier
identifier , *parameter-list*

The function-body has the form

function-body:
declaration-list *compound-statement*

The identifiers in the parameter list, and only those identifiers, may be declared in the declaration list. Any identifiers whose type is not given are taken to be **int**. The only storage class which may be specified is **register**; if it is specified, the corresponding actual parameter will be copied, if possible, into a register at the outset of the function.

A simple example of a complete function definition is

```
int max (a, b, c)
int a, b, c;
{
    int m;
    m = (a > b) ? a : b;
    return (m > c ? m : c);
}
```

Here 'int' is the type-specifier; 'max(a, b, c)' is the function-declarator; 'int a, b, c;' is the declaration-list for the formal parameters; '{ ... }' is the block giving the code for the statement. The parentheses in the return are not required.

C converts all float actual parameters to **double**, so formal parameters declared float have their declaration adjusted to read **double**. Also, since a reference to an array in any context (in particular as an actual parameter) is taken to mean a pointer to the first element of the array, declarations of formal parameters declared 'array of ...' are adjusted to read 'pointer to ...'. Finally, because neither structures nor functions can be passed to a function, it is useless to declare a formal parameter to be a structure or function (pointers to structures or functions are of course permitted).

A free return statement is supplied at the end of each function definition, so running off the end causes control, but no value, to be returned to the caller.

10.2 External data definitions

An external data definition has the form

data-definition:
declaration

The storage class of such data may be **extern** (which is the default) or **static**, but not **auto** or **register**.

11. Scope rules

A C program need not all be compiled at the same time: the source text of the program may be kept in several files, and precompiled routines may be loaded from libraries. Communication among the functions of a program may be carried out both through explicit calls and through manipulation of external data.

Therefore, there are two kinds of scope to consider: first, what may be called the *lexical scope* of an identifier, which is essentially the region of a program during which it may be used without drawing 'undefined identifier' diagnostics; and second, the scope associated with external identifiers, which is characterized by the rule that references to the same external identifier are references to the same object.

11.1 Lexical scope

The lexical scope of identifiers declared in external definitions persists from the definition through the end of the file in which they appear. The lexical scope of identifiers which are formal parameters persists through the function with which they are associated. The lexical scope of identifiers declared at the head of blocks persists until the end of the block. The lexical scope of labels is the whole of the function in which they appear.

Because all references to the same external identifier refer to the same object (see §11.2) the compiler checks all declarations of the same external identifier for compatibility; in effect their scope is increased to the whole file in which they appear.

In all cases, however, if an identifier is explicitly declared at the head of a block, including the block constituting a function, any declaration of that identifier outside the block is suspended until the end of the block.

Remember also (§8.5) that identifiers associated with ordinary variables on the one hand and those associated with structure and union members and tags on the other form two disjoint classes which do not conflict. `typedef` names are in the same class as ordinary identifiers. They may be redeclared in inner blocks, but an explicit type must be given in the inner declaration:

```
typedef float distance;  
...  
{ auto int distance;  
...  
}
```

The `int` must be present in the second declaration, or it would be taken to be a declaration with no declarators and type `distance`.*

11.2 Scope of externals

If a function declares an identifier to be `extern`, then somewhere among the files or libraries constituting the complete program there must be an external definition for the identifier. All functions in a given program which refer to the same external identifier refer to the same object, so care must be taken that the type and extent specified in the definition are compatible with those specified by each function which references the data.

In PDP-11 C, compatible external definitions of the same identifier may be present in several of the separately-compiled pieces of a complete program, or even twice within the same program file, with the limitation that the identifier may be initialized in at most one of the definitions. In other operating systems, however, the compiler must know in just which file the storage for the identifier is allocated, and in which file the identifier is merely being referred to. The appearance of the `extern` keyword in an external definition indicates that storage for the identifiers being declared will be allocated in another file. Thus in a multi-file program, an

*It is agreed that the ice is thin here.

external data definition without the `extern` specifier must appear in exactly one of the files. Any other files which wish to give an external definition for the identifier must include the `extern` in the definition. The identifier can be initialized only in the declaration where storage is allocated.

Identifiers declared `static` at the top level in external definitions are not visible in other files.

12. Compiler control lines

The C compiler contains a preprocessor capable of macro substitution, conditional compilation, and inclusion of named files. Lines beginning with `#` communicate with this preprocessor. These lines have syntax independent of the rest of the language; they may appear anywhere and have effect which lasts (independent of scope) until the end of the source program file.

12.1 Token replacement

A compiler-control line of the form

```
# define identifier token-string
```

(note: no trailing semicolon) causes the preprocessor to replace subsequent instances of the identifier with the given string of tokens. A line of the form

```
# define identifier( identifier , ... , identifier ) token-string
```

where there is no space between the first identifier and the `(`, is a macro definition with arguments. Subsequent instances of the first identifier followed by a `(`, a sequence of tokens delimited by commas, and a `)` are replaced by the token string in the definition. Each occurrence of an identifier mentioned in the formal parameter list of the definition is replaced by the corresponding token string from the call. The actual arguments in the call are token strings separated by commas; however commas in quoted strings or protected by parentheses do not separate arguments. The number of formal and actual parameters must be the same. Text inside a string or a character constant is not subject to replacement.

In both forms the replacement string is rescanned for more defined identifiers. In both forms a long definition may be continued on another line by writing `\` at the end of the line to be continued.

This facility is most valuable for definition of 'manifest constants', as in

```
# define TABSIZE 100
...
int table[TABSIZE];
```

A control line of the form

```
# undef identifier
```

causes the identifier's preprocessor definition to be forgotten.

12.2 File inclusion

A compiler control line of the form

```
# include "filename"
```

causes the replacement of that line by the entire contents of the file *filename*.

The named file is searched for first in the directory of the original source file, and then in a sequence of standard places. Alternatively, a control line of the form

```
# include <filename>
```

searches only the standard places, and not the directory of the source file.

Includes may be nested.

12.3 Conditional compilation

A compiler control line of the form

if constant-expression

checks whether the constant expression (see §15) evaluates to non-zero. A control line of the form

ifdef identifier

checks whether the identifier is currently defined in the preprocessor; that is, whether it has been the subject of a *#define* control line. A control line of the form

ifndef identifier

checks whether the identifier is currently undefined in the preprocessor.

All three forms are followed by an arbitrary number of lines, possibly containing a control line

else

and then by a control line

endif

If the checked condition is true then any lines between *#else* and *#endif* are ignored. If the checked condition is false then any lines between the test and an *#else* or, lacking an *#else*, the *#endif*, are ignored.

These constructions may be nested.

12.4 Line control

For the benefit of other preprocessors which generate C programs, a line of the form

line constant identifier

causes the compiler to believe, for purposes of error diagnostics, that the next line number is given by the constant and the current input file is named by the identifier. If the identifier is absent the remembered file name does not change.

13. Implicit declarations

It is not always necessary to specify both the storage class and the type of identifiers in a declaration. Sometimes the storage class is supplied by the context: in external definitions, and in declarations of formal parameters and structure members. In a declaration inside a function, if a storage class but no type is given, the identifier is assumed to be *int*; if a type but no storage class is indicated, the identifier is assumed to be *auto*. An exception to the latter rule is made for functions, since *auto* functions are meaningless (C being incapable of compiling code into the stack). If the type of an identifier is 'function returning ...', it is implicitly declared to be *extern*.

In an expression, an identifier followed by (and not currently declared is contextually declared to be 'function returning *int*'.

14. Types revisited

This section summarizes the operations which can be performed on objects of certain types.

14.1 Structures and unions

There are only two things that can be done with a structure or union: name one of its members (by means of the *.* operator); or take its address (by unary *&*). Other operations, such as assigning from or to it or passing it as a parameter, draw an error message. In the future, it is expected that these operations, but not necessarily others, will be allowed.

§7.1 says that in a direct or indirect structure reference (with *.* or *->*) the name on the

right must be a member of the structure named or pointed to by the expression on the left. To allow an escape from the typing rules, this restriction is not firmly enforced by the compiler. In fact, any lvalue is allowed before '.', and that lvalue is then assumed to have the form of the structure of which the name on the right is a member. Also, the expression before a '->' is required only to be a pointer or an integer. If a pointer, it is assumed to point to a structure of which the name on the right is a member. If an integer, it is taken to be the absolute address, in machine storage units, of the appropriate structure.

Such constructions are non-portable.

14.2 Functions

There are only two things that can be done with a function: call it, or take its address. If the name of a function appears in an expression not in the function-name position of a call, a pointer to the function is generated. Thus, to pass one function to another, one might say

```
int f();
...
g(f);
```

Then the definition of *g* might read

```
g(funcp)
int (*funcp)();
{
    ...
    (*funcp)();
    ...
}
```

Notice that *f* was declared explicitly in the calling routine since its first appearance was not followed by (.

14.3 Arrays, pointers, and subscripting

Every time an identifier of array type appears in an expression, it is converted into a pointer to the first member of the array. Because of this conversion, arrays are not lvalues. By definition, the subscript operator [] is interpreted in such a way that 'E1[E2]' is identical to '*((E1) + (E2))'. Because of the conversion rules which apply to +, if E1 is an array and E2 an integer, then E1[E2] refers to the E2-th member of E1. Therefore, despite its asymmetric appearance, subscripting is a commutative operation.

A consistent rule is followed in the case of multi-dimensional arrays. If E is an *n*-dimensional array of rank $i \times j \times \dots \times k$, then E appearing in an expression is converted to a pointer to an (*n*-1)-dimensional array with rank $j \times \dots \times k$. If the * operator, either explicitly or implicitly as a result of subscripting, is applied to this pointer, the result is the pointed-to (*n*-1)-dimensional array, which itself is immediately converted into a pointer.

For example, consider

```
int x[3][5];
```

Here *x* is a 3×5 array of integers. When *x* appears in an expression, it is converted to a pointer to (the first of three) 5-membered arrays of integers. In the expression 'x[i]', which is equivalent to '* (x+i)', *x* is first converted to a pointer as described; then *i* is converted to the type of *x*, which involves multiplying *i* by the length the object to which the pointer points, namely 5 integer objects. The results are added and indirection applied to yield an array (of 5 integers) which in turn is converted to a pointer to the first of the integers. If there is another subscript the same argument applies again; this time the result is an integer.

It follows from all this that arrays in C are stored row-wise (last subscript varies fastest) and that the first subscript in the declaration helps determine the amount of storage consumed

by an array but plays no other part in subscript calculations.

15. Constant expressions

In several places C requires expressions which evaluate to a constant: after `case`, as array bounds, and in initializers. In the first two cases, the expression can involve only integer constants, character constants, and `sizeof` expressions, possibly connected by the binary operators

+ - * / % & | ^ << >> == != < > <= >=

or by the unary operators

- ~

or by the ternary operator

? :

Parentheses can be used for grouping, but not for function calls.

A bit more latitude is permitted for initializers; besides constant expressions as discussed above, one can also apply the unary `&` operator to external or static objects, and to external or static arrays subscripted with a constant expression. The unary `&` can also be applied implicitly by appearance of unsubscripted arrays and functions. The basic rule is that initializers must evaluate either to a constant or to the address of a previously declared external or static object plus or minus a constant.

16. Grammar revisited.

This section repeats the grammar of C in notation somewhat different than given before. The description below is adapted directly from a YACC grammar actually used by several compilers; thus it may (aside from possible editing errors) be regarded as authentic. The notation is pure YACC with the exception that the '!' separating alternatives for a production is omitted, since alternatives are always on separate lines; the ';' separating productions is omitted since a blank line is left between productions.

The lines with '%term' name the terminal symbols, which are either commented upon or should be self-evident. The lines with '%left,' '%right,' and '%binary' indicate whether the listed terminals are left-associative, right-associative, or non-associative, and describe a precedence structure. The precedence (binding strength) increases as one reads down the page. When the construction '%prec x' appears the precedence of the rule is that of the terminal x; otherwise the precedence of the rule is that of its leftmost terminal.

```
%term NAME
%term STRING
%term ICON
%term FCON
%term PLUS
%term MINUS
%term MUL
%term AND
%term QUEST
%term COLON
%term ANDAND
%term OROR
%term ASOP      /* old-style = + etc. */
%term RELOP     /* < = > = < > */
%term EQUOP     /* == != */
%term DIVOP     /* / % */
%term OR        /* | */
%term EXOR      /* ^ */
%term SHIFTOP   /* << >> */
%term INCOP     /* ++ -- */
%term UNOP      /* ! ~ */
%term STROP     /* . - > */

%term TYPE      /* int, char, long, float, double, unsigned, short */
%term CLASS     /* extern, register, auto, static, typedef */
%term STRUCT    /* struct or union */
%term RETURN
%term GOTO
%term IF
%term ELSE
%term SWITCH
%term BREAK
%term CONTINUE
%term WHILE
%term DO
%term FOR
%term DEFAULT
%term CASE
%term SIZEOF
```

%term LP /* (*/
%term RP /*) */
%term LC /* { */
%term RC /* } */
%term LB /* [*/
%term RB /*] */
%term CM /* , */
%term SM /* ; */
%term ASSIGN /* = */

%left CM
%right ASOP ASSIGN
%right QUEST COLON
%left OROR
%left ANDAND
%left OROP
%left AND
%binary EQUOP
%binary RELOP
%left SHIFTOP
%left PLUS MINUS
%left MUL DIVOP
%right UNOP
%right INCOP SIZEOF
%left LB LP STROP

program: ext_def_list

ext_def_list: ext_def_list external_def
 /* empty */

external_def: optattrib SM
 optattrib init_dcl_list SM
 optattrib fdeclarator function_body

function_body: dcl_list compoundstmt

dcl_list: dcl_list declaration
 /* empty */

declaration: specifiers declarator_list SM
 specifiers SM

optattrib: specifiers
 /* empty */

specifiers: CLASS type
 type CLASS
 CLASS
 type

type: TYPE
TYPE TYPE
struct_dcl

struct_dcl: STRUCT NAME LC type_dcl_list RC
STRUCT LC type_dcl_list RC
STRUCT NAME

type_dcl_list: type_declaration
type_dcl_list type_declaration

type_declaration: type_declarator_list SM
struct_dcl SM
type SM

declarator_list: declarator
declarator_list CM declarator

declarator: fdeclarator
nfdeclarator
nfdeclarator COLON con_e %prec CM
COLON con_e %prec CM

nfdeclarator: MUL nfdeclarator
nfdeclarator LP RP
nfdeclarator LB RB
nfdeclarator LB con_e RB
NAME
LP nfdeclarator RP

fdeclarator: MUL fdeclarator
fdeclarator LP RP
fdeclarator LB RB
fdeclarator LB con_e RB
LP fdeclarator RP
NAME LP name_list RP
NAME LP RP

name_list: NAME
name_list CM NAME

init_dcl_list: init_declarator %prec CM
init_dcl_list CM init_declarator

init_declarator: nfdeclarator
nfdeclarator ASSIGN initializer
nfdeclarator initializer
fdeclarator

init_list: initializer %prec CM
init_list CM initializer

initializer: e %prec CM
LC init_list RC

LC init_list CM RC

compoundstmt: LC del_list stmt_list RC

stmt_list: stmt_list statement
/* empty */

statement: e SM
compoundstmt
IF LP e RP statement
IF LP e RP statement ELSE statement
WHILE LP e RP statement
DO statement WHILE LP e RP SM
FOR LP opt_e SM opt_e SM opt_e RP statement
SWITCH LP e RP statement
BREAK SM
CONTINUE SM
RETURN SM
RETURN e SM
GOTO NAME SM
SM
label statement

label: NAME COLON
CASE con_e COLON
DEFAULT COLON

con_e: e %prec CM

opt_e: e
/* empty */

elist: e %prec CM
elist CM e


e: e MUL e
e CM e
e DIVOP e
e PLUS e
e MINUS e
e SHIFTOP e
e RELOP e
e EQUOP e
e AND e
e OROP e
e ANDAND e
e OROR e
e MUL ASSIGN e
e DIVOP ASSIGN e
e PLUS ASSIGN e
e MINUS ASSIGN e
e SHIFTOP ASSIGN e
e AND ASSIGN e
e OROP ASSIGN e


e QUEST e COLON e
e ASOP e
e ASSIGN e
term

term: term INCOP
MUL term
AND term
MINUS term
UNOP term
INCOP term
SIZEOF term
LP type_name RP term %prec STROP
SIZEOF LP type_name RP %prec SIZEOF
term LB e RB
term LP RP
term LP elist RP
term STROP NAME
NAME
ICON
FCON
STRING
LP e RP

type_name: type abst_decl

abst_decl: /* empty */
LP RP
LP abst_decl RP LP RP
MUL abst_decl
abst_decl LB RB
abst_decl LB con_e RB
LP abst_decl RP





Programming in C — A Tutorial

Brian W. Kernighan

Bell Laboratories, Murray Hill, N. J.

1. Introduction

C is a computer language available on the GCOS and UNIX operating systems at Murray Hill and (in preliminary form) on OS/360 at Holmdel. C lets you write your programs clearly and simply — it has decent control flow facilities so your code can be read straight down the page, without labels or GOTO's; it lets you write code that is compact without being too cryptic; it encourages modularity and good program organization; and it provides good data-structuring facilities.

This memorandum is a tutorial to make learning C as painless as possible. The first part concentrates on the central features of C; the second part discusses those parts of the language which are useful (usually for getting more efficient and smaller code) but which are not necessary for the new user. This is *not* a reference manual. Details and special cases will be skipped ruthlessly, and no attempt will be made to cover every language feature. The order of presentation is hopefully pedagogical instead of logical. Users who would like the full story should consult the *C Reference Manual* by D. M. Ritchie [1], which should be read for details anyway. Runtime support is described in [2] and [3]; you will have to read one of these to learn how to compile and run a C program.

We will assume that you are familiar with the mysteries of creating files, text editing, and the like in the operating system you run on, and that you have programmed in some language before.

2. A Simple C Program

```
main( ) {  
    printf("hello, world");  
}
```

A C program consists of one or more *functions*, which are similar to the functions and subroutines of a Fortran program or the procedures of PL/I, and perhaps some external data definitions. **main** is such a function, and in fact all C programs must have a **main**. Execution of the program begins at the first statement of **main**. **main** will usually invoke other functions to perform its job, some coming from the same program, and others from libraries.

One method of communicating data between functions is by arguments. The parentheses following the function name surround the argument list; here **main** is a function of no arguments, indicated by (). The {} enclose the statements of the function. Individual statements end with a semicolon but are otherwise free-format.

`printf` is a library function which will format and print output on the terminal (unless some other destination is specified). In this case it prints

```
hello, world
```

A function is invoked by naming it, followed by a list of arguments in parentheses. There is no `CALL` statement as in Fortran or PL/I.

3. A Working C Program; Variables; Types and Type Declarations

Here's a bigger program that adds three integers and prints their sum.

```
main( ) {
    int a, b, c, sum;
    a = 1; b = 2; c = 3;
    sum = a + b + c;
    printf("sum is %d", sum);
}
```

Arithmetic and the assignment statements are much the same as in Fortran (except for the semicolons) or PL/I. The format of C programs is quite free. We can put several statements on a line if we want, or we can split a statement among several lines if it seems desirable. The split may be between any of the operators or variables, but *not* in the middle of a name or operator. As a matter of style, spaces, tabs, and newlines should be used freely to enhance readability.

C has four fundamental *types* of variables:

```
int   integer (PDP-11: 16 bits; H6070: 36 bits; IBM360: 32 bits)
char  one byte character (PDP-11, IBM360: 8 bits; H6070: 9 bits)
float single-precision floating point
double double-precision floating point
```

There are also *arrays* and *structures* of these basic types, *pointers* to them and *functions* that return them, all of which we will meet shortly.

All variables in a C program must be declared, although this can sometimes be done implicitly by context. Declarations must precede executable statements. The declaration

```
int a, b, c, sum;
```

declares `a`, `b`, `c`, and `sum` to be integers.

Variable names have one to eight characters, chosen from A-Z, a-z, 0-9, and `_`, and start with a non-digit. Stylistically, it's much better to use only a single case and give functions and external variables names that are unique in the first six characters. (Function and external variable names are used by various assemblers, some of which are limited in the size and case of identifiers they can handle.) Furthermore, keywords and library functions may only be recognized in one case.

4. Constants

We have already seen decimal integer constants in the previous example — 1, 2, and 3. Since C is often used for system programming and bit-manipulation, octal numbers are an important part of the language. In C, any number that begins with 0 (zero!) is an octal integer (and hence can't have any 8's or 9's in it). Thus 0777 is an octal constant, with decimal value 511.

A "character" is one byte (an inherently machine-dependent concept). Most often this is expressed as a *character constant*, which is one character enclosed in single quotes. However, it may be any quantity that fits in a byte, as in `flags` below:

```

char quest, newline, flags;
quest = '?';
newline = '\n';
flags = 077;

```

The sequence `'\n'` is C notation for "newline character", which, when printed, skips the terminal to the beginning of the next line. Notice that `'\n'` represents only a single character. There are several other "escapes" like `'\n'` for representing hard-to-get or invisible characters, such as `'\t'` for tab, `'\b'` for backspace, `'\0'` for end of file, and `'\'` for the backslash itself.

`float` and `double` constants are discussed in section 26.

5. Simple I/O – `getchar`, `putchar`, `printf`

```

main() {
    char c;
    c = getchar();
    putchar(c);
}

```

`getchar` and `putchar` are the basic I/O library functions in C. `getchar` fetches one character from the standard input (usually the terminal) each time it is called, and returns that character as the value of the function. When it reaches the end of whatever file it is reading, thereafter it returns the character represented by `'\0'` (ascii NUL, which has value zero). We will see how to use this very shortly.

`putchar` puts one character out on the standard output (usually the terminal) each time it is called. So the program above reads one character and writes it back out. By itself, this isn't very interesting, but observe that if we put a loop around this, and add a test for end of file, we have a complete program for copying one file to another.

`printf` is a more complicated function for producing formatted output. We will talk about only the simplest use of it. Basically, `printf` uses its first argument as formatting information, and any successive arguments as variables to be output. Thus

```
printf("hello, world\n");
```

is the simplest use — the string "hello, world\n" is printed out. No formatting information, no variables, so the string is dumped out verbatim. The newline is necessary to put this out on a line by itself. (The construction

```
"hello, world\n"
```

is really an array of `chars`. More about this shortly.)

More complicated, if `sum` is 6,

```
printf("sum is %d\n", sum);
```

prints

```
sum is 6
```

Within the first argument of `printf`, the characters `"%d"` signify that the next argument in the argument list is to be printed as a base 10 number.

Other useful formatting commands are `"%c"` to print out a single character, `"%s"` to print out an entire string, and `"%o"` to print a number as octal instead of decimal (no leading zero). For example,

```

n = 511;
printf("What is the value of %d in octal?", n);

```

```
printf (" %s! %d decimal is %o octal\n", "Right", n, n);
```

prints

What is the value of 511 in octal? Right! 511 decimal is 777 octal

Notice that there is no newline at the end of the first output line. Successive calls to `printf` (and/or `putchar`, for that matter) simply put out characters. No newlines are printed unless you ask for them. Similarly, on input, characters are read one at a time as you ask for them. Each line is generally terminated by a newline (`\n`), but there is otherwise no concept of record.

6. If; relational operators; compound statements

The basic conditional-testing statement in C is the if statement:

```
c = getchar( );
if ( c == '?' )
    printf("why did you type a question mark?\n");
```

The simplest form of if is

if (expression) statement

The condition to be tested is any expression enclosed in parentheses. It is followed by a statement. The expression is evaluated, and if its value is non-zero, the statement is executed. There's an optional `else` clause, to be described soon.

The character sequence `'=='` is one of the relational operators in C; here is the complete set:

```
-- equal to (.EQ. to Fortraners)
!= not equal to
> greater than
< less than
>= greater than or equal to
<= less than or equal to
```

The value of "expression relation expression" is 1 if the relation is true, and 0 if false. Don't forget that the equality test is `'=='`; a single `'='` causes an assignment, not a test, and invariably leads to disaster.

Tests can be combined with the operators `'&&'` (AND), `'||'` (OR), and `'!'` (NOT). For example, we can test whether a character is blank or tab or newline with

```
if ( c == ' ' || c == '\t' || c == '\n' ) ...
```

C guarantees that `'&&'` and `'||'` are evaluated left to right — we shall soon see cases where this matters.

One of the nice things about C is that the **statement** part of an if can be made arbitrarily complicated by enclosing a set of statements in `{}`. As a simple example, suppose we want to ensure that `a` is bigger than `b`, as part of a sort routine. The interchange of `a` and `b` takes three statements in C, grouped together by `{}`:

```
if ( a < b ) {
    t = a;
    a = b;
    b = t;
}
```

As a general rule in C, anywhere you can use a simple statement, you can use any compound statement, which is just a number of simple or compound ones enclosed in {}. There is no semicolon after the } of a compound statement, but there *is* a semicolon after the last non-compound statement inside the {}.

The ability to replace single statements by complex ones at will is one feature that makes C much more pleasant to use than Fortran. Logic (like the exchange in the previous example) which would require several GOTO's and labels in Fortran can and should be done in C without any, using compound statements.

7. While Statement; Assignment within an Expression; Null Statement

The basic looping mechanism in C is the **while** statement. Here's a program that copies its input to its output a character at a time. Remember that '\0' marks the end of file.

```
main( ) {
    char c;
    while( (c=getchar( )) != '\0' )
        putchar(c);
}
```

The **while** statement is a loop, whose general form is

```
while (expression) statement
```

Its meaning is

- (a) evaluate the expression
- (b) if its value is true (i.e., not zero)
 - do the statement, and go back to (a)

Because the expression is tested before the statement is executed, the statement part can be executed zero times, which is often desirable. As in the if statement, the expression and the statement can both be arbitrarily complicated, although we haven't seen that yet. Our example gets the character, assigns it to c, and then tests if it's a '\0'. If it is not a '\0', the statement part of the **while** is executed, printing the character. The **while** then repeats. When the input character is finally a '\0', the **while** terminates, and so does **main**.

Notice that we used an assignment statement

```
c = getchar( )
```

within an expression. This is a handy notational shortcut which often produces clearer code. (In fact it is often the only way to write the code cleanly. As an exercise, re-write the file-copy without using an assignment inside an expression.) It works because an assignment statement has a value, just as any other expression does. Its value is the value of the right hand side. This also implies that we can use multiple assignments like

```
x = y = z = 0;
```

Evaluation goes from right to left.

By the way, the extra parentheses in the assignment statement within the conditional were really necessary: if we had said

```
c = getchar( ) != '\0'
```

c would be set to 0 or 1 depending on whether the character fetched was an end of file or not. This is because in the absence of parentheses the assignment operator '=' is evaluated after the relational operator '!='. When in doubt, or even if not, parenthesize.

Since `putchar(c)` returns `c` as its function value, we could also copy the input to the output by nesting the calls to `getchar` and `putchar`:

```
main() {
    while( putchar(getchar( )) != '\0' );
}
```

What statement is being repeated? None, or technically, the *null* statement, because all the work is really done within the test part of the `while`. This version is slightly different from the previous one, because the final `\0` is copied to the output before we decide to stop.

8. Arithmetic

The arithmetic operators are the usual `+`, `-`, `*`, and `/` (truncating integer division if the operands are both `int`), and the remainder or mod operator `%`:

```
x = a%b;
```

sets `x` to the remainder after `a` is divided by `b` (i.e., `a mod b`). The results are machine dependent unless `a` and `b` are both positive.

In arithmetic, `char` variables can usually be treated like `int` variables. Arithmetic on characters is quite legal, and often makes sense:

```
c = c + 'A' - 'a';
```

converts a single lower case ascii character stored in `c` to upper case, making use of the fact that corresponding ascii letters are a fixed distance apart. The rule governing this arithmetic is that all `chars` are converted to `int` before the arithmetic is done. Beware that conversion may involve sign-extension — if the leftmost bit of a character is 1, the resulting integer might be negative. (This doesn't happen with genuine characters on any current machine.)

So to convert a file into lower case:

```
main() {
    char c;
    while( (c=getchar( )) != '\0' )
        if( 'A' <= c && c <= 'Z' )
            putchar(c+'a'-'A');
        else
            putchar(c);
}
```

Characters have different sizes on different machines. Further, this code won't work on an IBM machine, because the letters in the ebcidic alphabet are not contiguous.

9. Else Clause; Conditional Expressions

We just used an `else` after an `if`. The most general form of `if` is

```
if (expression) statement1 else statement2
```

the `else` part is optional, but often useful. The canonical example sets `x` to the minimum of `a` and `b`:

```
if (a < b)
    x = a;
else
    x = b;
```

Observe that there is a semicolon after `x = a`.

C provides an alternate form of conditional which is often more concise. It is called the "conditional expression" because it is a conditional which actually has a value and can be used anywhere an expression can. The value of

```
a < b ? a : b;
```

is **a** if **a** is less than **b**; it is **b** otherwise. In general, the form

```
expr1 ? expr2 : expr3
```

means "evaluate **expr1**. If it is not zero, the value of the whole thing is **expr2**; otherwise the value is **expr3**."

To set **x** to the minimum of **a** and **b**, then:

```
x = (a < b ? a : b);
```

The parentheses aren't necessary because '?' is evaluated before '=', but safety first.

Going a step further, we could write the loop in the lower-case program as

```
while( (c=getchar( )) != '\0' )
    putchar( ('A' <= c && c <= 'Z') ? c - 'A' + 'a' : c );
```

If's and else's can be used to construct logic that branches one of several ways and then rejoins, a common programming structure, in this way:

```
if(...)
    {...}
else if(...)
    {...}
else if(...)
    {...}
else
    {...}
```

The conditions are tested in order, and exactly one block is executed — either the first one whose **if** is satisfied, or the one for the last **else**. When this block is finished, the next statement executed is the one after the last **else**. If no action is to be taken for the "default" case, omit the last **else**.

For example, to count letters, digits and others in a file, we could write

```
main( ) {
    int let, dig, other, c;
    let = dig = other = 0;
    while( (c=getchar( )) != '\0' )
        if( ('A' <= c && c <= 'Z') || ('a' <= c && c <= 'z') ) ++let;
        else if( '0' <= c && c <= '9' ) ++dig;
        else ++other;
    printf("%d letters, %d digits, %d others\n", let, dig, other);
}
```

The '++' operator means "increment by 1"; we will get to it in the next section.

10. Increment and Decrement Operators

In addition to the usual '-', C also has two other interesting unary operators, '++' (increment) and '--' (decrement). Suppose we want to count the lines in a file.

```
main( ) {
    int c,n;
    n = 0;
```

```

while( (c=getchar( )) != '\0' )
    if( c == '\n' )
        ++n;
printf("%d lines\n", n);
}

```

`++n` is equivalent to `n=n+1` but clearer, particularly when `n` is a complicated expression. `++` and `--` can be applied only to `int`'s and `char`'s (and `pointers` which we haven't got to yet).

The unusual feature of `++` and `--` is that they can be used either before or after a variable. The value of `++k` is the value of `k` after it has been incremented. The value of `k++` is `k` before it is incremented. Suppose `k` is 5. Then

```
x = ++k;
```

increments `k` to 6 and then sets `x` to the resulting value, i.e., to 6. But

```
x = k++;
```

first sets `x` to 5, and then increments `k` to 6. The incrementing effect of `++k` and `k++` is the same, but their values are respectively 5 and 6. We shall soon see examples where both of these uses are important.

11. Arrays

In C, as in Fortran or PL/I, it is possible to make arrays whose elements are basic types. Thus we can make an array of 10 integers with the declaration

```
int x[10];
```

The square brackets mean *subscripting*; parentheses are used only for function references. Array indexes begin at zero, so the elements of `x` are

```
x[0], x[1], x[2], ..., x[9]
```

If an array has `n` elements, the largest subscript is `n-1`.

Multiple-dimension arrays are provided, though not much used above two dimensions. The declaration and use look like

```
int name[10][20];
n = name[i+j][1] + name[k][2];
```

Subscripts can be arbitrary integer expressions. Multi-dimension arrays are stored by row (opposite to Fortran), so the rightmost subscript varies fastest; `name` has 10 rows and 20 columns.

Here is a program which reads a line, stores it in a buffer, and prints its length (excluding the newline at the end).

```

main( ) {
    int n, c;
    char line[100];
    n = 0;
    while( (c=getchar( )) != '\n' ) {
        if( n < 100 )
            line[n] = c;
        n++;
    }
    printf("length = %d\n", n);
}

```


As a more complicated problem, suppose we want to print the count for each line in the input, still storing the first 100 characters of each line. Try it as an exercise before looking at the solution:

```
main( ) {
    int n, c; char line[100];
    n = 0;
    while( (c=getchar( )) != '\0' )
        if( c == '\n' ) {
            printf("%d\n", n);
            n = 0;
        }
        else {
            if( n < 100 ) line[n] = c;
            n++;
        }
}
```

12. Character Arrays; Strings

Text is usually kept as an array of characters, as we did with `line[]` in the example above. By convention in C, the last character in a character array should be a `'\0'` because most programs that manipulate character arrays expect it. For example, `printf` uses the `'\0'` to detect the end of a character array when printing it out with a `'%s'`.

We can copy a character array `s` into another `t` like this:

```
i = 0;
while( (t[i]=s[i]) != '\0' )
    i++;
```

Most of the time we have to put in our own `'\0'` at the end of a string; if we want to print the line with `printf`, it's necessary. This code prints the character count before the line:

```
main( ) {
    int n;
    char line[100];
    n = 0;
    while( (line[n++] = getchar( )) != '\n' );
    line[n] = '\0';
    printf("%d:\t%s", n, line);
}
```

Here we increment `n` in the subscript itself, but only after the previous value has been used. The character is read, placed in `line[n]`, and only then `n` is incremented.

There is one place and one place only where C puts in the `'\0'` at the end of a character array for you, and that is in the construction

"stuff between double quotes"

The compiler puts a `'\0'` at the end automatically. Text enclosed in double quotes is called a *string*; its properties are precisely those of an (initialized) array of characters.

13. For Statement

The **for** statement is a somewhat generalized **while** that lets us put the initialization and increment parts of a loop into a single statement along with the test. The general form of the **for** is

```
for( initialization; expression; increment )
    statement
```

The meaning is exactly

```
initialization;
while( expression ) {
    statement
    increment;
}
```

Thus, the following code does the same array copy as the example in the previous section:

```
for( i=0; (t[i]=s[i]) != '\0'; i++ );
```

This slightly more ornate example adds up the elements of an array:

```
sum = 0;
for( i=0; i<n; i++ )
    sum = sum + array[i];
```

In the **for** statement, the initialization can be left out if you want, but the semicolon has to be there. The increment is also optional. It is *not* followed by a semicolon. The second clause, the test, works the same way as in the **while**: if the expression is true (not zero) do another loop, otherwise get on with the next statement. As with the **while**, the **for** loop may be done zero times. If the expression is left out, it is taken to be always true, so

```
for( ; ; ) ...
```

and

```
while( 1 ) ...
```

are both infinite loops.

You might ask why we use a **for** since it's so much like a **while**. (You might also ask why we use a **while** because...) The **for** is usually preferable because it keeps the code where it's used and sometimes eliminates the need for compound statements, as in this code that zeros a two-dimensional array:

```
for( i=0; i<n; i++ )
    for( j=0; j<m; j++ )
        array[i][j] = 0;
```

14. Functions; Comments

Suppose we want, as part of a larger program, to count the occurrences of the **ascii** characters in some input text. Let us also map illegal characters (those with value >127 or <0) into one pile. Since this is presumably an isolated part of the program, good practice dictates making it a separate function. Here is one way:

```

main( ) {
    int hist[129];          /* 128 legal chars + 1 illegal group */
    ...
    count(hist, 128);      /* count the letters into hist */
    printf( ... );        /* comments look like this; use them */
    ...                    /* anywhere blanks, tabs or newlines could appear */
}

count(buf, size)
int size, buf ]; {
    int i, c;
    for( i=0; i<=size; i++ )
        buf[i] = 0;          /* set buf to zero */
    while( (c=getchar( )) != '\0' ) { /* read til eof */
        if( c > size || c < 0 )
            c = size;        /* fix illegal input */
        buf[c]++;
    }
    return;
}

```

We have already seen many examples of calling a function, so let us concentrate on how to *define* one. Since `count` has two arguments, we need to declare them, as shown, giving their types, and in the case of `buf`, the fact that it is an array. The declarations of arguments go *between* the argument list and the opening '{'. There is no need to specify the size of the array `buf`, for it is defined outside of `count`.

The `return` statement simply says to go back to the calling routine. In fact, we could have omitted it, since a return is implied at the end of a function.

What if we wanted `count` to return a value, say the number of characters read? The `return` statement allows for this too:

```

int i, c, nchar;
nchar = 0;
...
while( (c=getchar( )) != '\0' ) {
    if( c > size || c < 0 )
        c = size;
    buf[c]++;
    nchar++;
}
return(nchar);

```

Any expression can appear within the parentheses. Here is a function to compute the minimum of two integers:

```

min(a, b)
int a, b; {
    return( a < b ? a : b );
}

```

To copy a character array, we could write the function

```

strcpy(s1, s2)      /* copies s1 to s2 */
  char s1[ ], s2[ ]; {
    int i;
    for( i = 0; (s2[i] = s1[i]) != '\0'; i++ );
  }

```

As is often the case, all the work is done by the assignment statement embedded in the test part of the `for`. Again, the declarations of the arguments `s1` and `s2` omit the sizes, because they don't matter to `strcpy`. (In the section on pointers, we will see a more efficient way to do a string copy.)

There is a subtlety in function usage which can trap the unsuspecting Fortran programmer. Simple variables (not arrays) are passed in C by "call by value", which means that the called function is given a copy of its arguments, and doesn't know their addresses. This makes it impossible to change the value of one of the actual input arguments.

There are two ways out of this dilemma. One is to make special arrangements to pass to the function the address of a variable instead of its value. The other is to make the variable a global or external variable, which is known to each function by its name. We will discuss both possibilities in the next few sections.

15. Local and External Variables

If we say

```

f() {
    int x;
    ...
}
g() {
    int x;
    ...
}

```

each `x` is *local* to its own routine — the `x` in `f` is unrelated to the `x` in `g`. (Local variables are also called "automatic".) Furthermore each local variable in a routine appears only when the function is called, and *disappears* when the function is exited. Local variables have no memory from one call to the next and must be explicitly initialized upon each entry. (There is a `static` storage class for making local variables with memory; we won't discuss it.)

As opposed to local variables, *external variables* are defined external to all functions, and are (potentially) available to all functions. External storage always remains in existence. To make variables external we have to *define* them external to all functions, and, wherever we want to use them, make a *declaration*.

```

main() {
    extern int nchar, hist[ ];
    ...
    count( );
    ...
}

```

```

count( ) {
    extern int nchar, hist[ ];
    int i, c;
    ...
}

int hist[129]; /* space for histogram */
int nchar; /* character count */

```

Roughly speaking, any function that wishes to access an external variable must contain an **extern** declaration for it. The declaration is the same as others, except for the added keyword **extern**. Furthermore, there must somewhere be a *definition* of the external variables external to all functions.

External variables can be initialized; they are set to zero if not explicitly initialized. In its simplest form, initialization is done by putting the value (which must be a constant) after the definition:

```

int nchar 0;
char flag 'f';
etc.

```

This is discussed further in a later section.

This ends our discussion of what might be called the central core of C. You now have enough to write quite substantial C programs, and it would probably be a good idea if you paused long enough to do so. The rest of this tutorial will describe some more ornate constructions, useful but not essential.

16. Pointers

A *pointer* in C is the address of something. It is a rare case indeed when we care what the specific address itself is, but pointers are a quite common way to get at the contents of something. The unary operator **&** is used to produce the address of an object, if it has one. Thus

```

int a, b;
b = &a;

```

puts the address of **a** into **b**. We can't do much with it except print it or pass it to some other routine, because we haven't given **b** the right kind of declaration. But if we declare that **b** is indeed a *pointer* to an integer, we're in good shape:

```

int a, *b, c;
b = &a;
c = *b;

```

b contains the address of **a** and '**c = *b**' means to use the value in **b** as an address, i.e., as a pointer. The effect is that we get back the contents of **a**, albeit rather indirectly. (It's always the case that '***&x**' is the same as **x** if **x** has an address.)

The most frequent use of pointers in C is for walking efficiently along arrays. In fact, in the implementation of an array, the array name represents the address of the zeroth element of the array, so you can't use it on the left side of an expression. (You can't change the address of something by assigning to it.) If we say

```

char *y;
char x[100];

```

y is of type pointer to character (although it doesn't yet point anywhere). We can make **y** point to an element of **x** by either of

point to an element of `x` by either of

```
y = &x[0];
y = x;
```

Since `x` is the address of `x[0]` this is legal and consistent.

Now `*y` gives `x[0]`. More importantly,

```
*(y+1) gives x[1]
*(y+i) gives x[i]
```

and the sequence

```
y = &x[0];
y++;
```

leaves `y` pointing at `x[1]`.

Let's use pointers in a function `length` that computes how long a character array is. Remember that by convention all character arrays are terminated with a `'\0'`. (And if they aren't, this program will blow up inevitably.) The old way:

```
length(s)
char s[ ]; {
    int n;
    for( n=0; s[n] != '\0'; )
        n++;
    return(n);
}
```

Rewriting with pointers gives

```
length(s)
char *s; {
    int n;
    for( n=0; *s != '\0'; s++ )
        n++;
    return(n);
}
```

You can now see why we have to say what kind of thing `s` points to — if we're to increment it with `s++` we have to increment it by the right amount.

The pointer version is more efficient (this is almost always true) but even more compact is

```
for( n=0; *s++ != '\0'; n++ );
```

The `*s` returns a character; the `++` increments the pointer so we'll get the next character next time around. As you can see, as we make things more efficient, we also make them less clear. But `*s++` is an idiom so common that you have to know it.

Going a step further, here's our function `strcpy` that copies a character array `s` to another `t`.

```
strcpy(s,t)
char *s, *t; {
    while(*t++ = *s++);
}
```

We have omitted the test against `'\0'`, because `'\0'` is identically zero; you will often see the code this way. (You *must* have a space after the `'='`: see section 25.)

For arguments to a function, and there only, the declarations

```
char s[ ];
char *s;
```

are equivalent — a pointer to a type, or an array of unspecified size of that type, are the same thing.

If this all seems mysterious, copy these forms until they become second nature. You don't often need anything more complicated.

17. Function Arguments

Look back at the function `strcpy` in the previous section. We passed it two string names as arguments, then proceeded to clobber both of them by incrementation. So how come we don't lose the original strings in the function that called `strcpy`?

As we said before, C is a "call by value" language: when you make a function call like `f(x)`, the *value* of `x` is passed, not its address. So there's no way to *alter* `x` from inside `f`. If `x` is an array (`char x[10]`) this isn't a problem, because `x` is an address anyway, and you're not trying to change it, just what it addresses. This is why `strcpy` works as it does. And it's convenient not to have to worry about making temporary copies of the input arguments.

But what if `x` is a scalar and you do want to change it? In that case, you have to pass the *address* of `x` to `f`, and then use it as a pointer. Thus for example, to interchange two integers, we must write

```
flip(x, y)
int *x, *y; {
    int temp;
    temp = *x;
    *x = *y;
    *y = temp;
}
```

and to call `flip`, we have to pass the addresses of the variables:

```
flip (&a, &b);
```

18. Multiple Levels of Pointers; Program Arguments

When a C program is called, the arguments on the command line are made available to the main program as an argument count `argc` and an array of character strings `argv` containing the arguments. Manipulating these arguments is one of the most common uses of multiple levels of pointers ("pointer to pointer to ..."). By convention, `argc` is greater than zero; the first argument (in `argv[0]`) is the command name itself.

Here is a program that simply echoes its arguments.

```
main(argc, argv)
int argc;
char **argv; {
    int i;
    for( i = 1; i < argc; i++ )
        printf("%s ", argv[i]);
    putchar('\n');
}
```

Step by step: `main` is called with two arguments, the argument count and the array of arguments. `argv` is a pointer to an array, whose individual elements are pointers to arrays of char-

acters. The zeroth argument is the name of the command itself, so we start to print with the first argument, until we've printed them all. Each `argv[i]` is a character array, so we use a `'%s'` in the `printf`.

You will sometimes see the declaration of `argv` written as

```
char *argv[ ];
```

which is equivalent. But we can't use `char argv[][]`, because both dimensions are variable and there would be no way to figure out how big the array is.

Here's a bigger example using `argc` and `argv`. A common convention in C programs is that if the first argument is `'-'`, it indicates a flag of some sort. For example, suppose we want a program to be callable as

```
prog -abc arg1 arg2 ...
```

where the `'-'` argument is optional; if it is present, it may be followed by any combination of a, b, and c.

```
main(argc, argv)
int argc;
char **argv; {
    ...
    aflag = bflag = cflag = 0;
    if( argc > 1 && argv[1][0] == '-' ) {
        for( i=1; (c=argv[1][i]) != '\0'; i++ )
            if( c=='a' )
                aflag++;
            else if( c=='b' )
                bflag++;
            else if( c=='c' )
                cflag++;
            else
                printf("%c?\n", c);
        --argc;
        ++argv;
    }
    ...
}
```

There are several things worth noticing about this code. First, there is a real need for the left-to-right evaluation that `&&` provides; we don't want to look at `argv[1]` unless we know it's there. Second, the statements

```
--argc;
++argv;
```

let us march along the argument list by one position, so we can skip over the flag argument as if it had never existed — the rest of the program is independent of whether or not there was a flag argument. This only works because `argv` is a pointer which can be incremented.

19. The Switch Statement; Break; Continue

The `switch` statement can be used to replace the multi-way test we used in the last example. When the tests are like this:

```
if( c == 'a' ) ...
else if( c == 'b' ) ...
else if( c == 'c' ) ...
else ...
```


testing a value against a series of *constants*, the switch statement is often clearer and usually gives better code. Use it like this:

```
switch( c ) {
    case 'a':
        aflag ++;
        break;
    case 'b':
        bflag ++;
        break;
    case 'c':
        cflag ++;
        break;
    default:
        printf("%c?\n", c);
        break;
}
```

The **case** statements label the various actions we want; **default** gets done if none of the other cases are satisfied. (A **default** is optional; if it isn't there, and none of the cases match, you just fall out the bottom.)

The **break** statement in this example is new. It is there because the cases are just labels, and after you do one of them, you *fall through* to the next unless you take some explicit action to escape. This is a mixed blessing. On the positive side, you can have multiple cases on a single statement; we might want to allow both upper and lower case letters in our flag field, so we could say

```
case 'a': case 'A': ...
case 'b': case 'B': ...
etc.
```

But what if we just want to get out after doing **case 'a'**? We could get out of a **case** of the **switch** with a label and a **goto**, but this is really ugly. The **break** statement lets us exit without either **goto** or label.

```
switch( c ) {
    case 'a':
        aflag ++;
        break;
    case 'b':
        bflag ++;
        break;
    ...
}
/* the break statements get us here directly */
```

The **break** statement also works in **for** and **while** statements — it causes an immediate exit from the loop.

The **continue** statement works *only* inside **for**'s and **while**'s; it causes the next iteration of the loop to be started. This means it goes to the increment part of the **for** and the test part of the **while**. We could have used a **continue** in our example to get on with the next iteration of the **for**, but it seems clearer to use **break** instead.

20. Structures

The main use of structures is to lump together collections of disparate variable types, so they can conveniently be treated as a unit. For example, if we were writing a compiler or assembler, we might need for each identifier information like its name (a character array), its source line number (an integer), some type information (a character, perhaps), and probably a usage count (another integer).

```
char  id[10];
int   line;
char  type;
int   usage;
```

We can make a structure out of this quite easily. We first tell C what the structure will look like, that is, what kinds of things it contains; after that we can actually reserve storage for it, either in the same statement or separately. The simplest thing is to define it and allocate storage all at once:

```
struct {
    char  id[10];
    int   line;
    char  type;
    int   usage;
} sym;
```

This defines `sym` to be a structure with the specified shape; `id`, `line`, `type` and `usage` are *members* of the structure. The way we refer to any particular member of the structure is

```
structure-name . member
```

as in

```
sym.type = 077;
if( sym.usage == 0 ) ...
while( sym.id[j++ ] ) ...
etc.
```

Although the names of structure members never stand alone, they still have to be unique — there can't be another `id` or `usage` in some other structure.

So far we haven't gained much. The advantages of structures start to come when we have arrays of structures, or when we want to pass complicated data layouts between functions. Suppose we wanted to make a symbol table for up to 100 identifiers. We could extend our definitions like

```
char  id[100][10];
int   line[100];
char  type[100];
int   usage[100];
```

but a structure lets us rearrange this spread-out information so all the data about a single identifier is collected into one lump:

```
struct {
    char  id[10];
    int   line;
    char  type;
    int   usage;
} sym[100];
```

This makes `sym` an array of structures; each array element has the specified shape. Now we can refer to members as

```
sym[i].usage++; /* increment usage of i-th identifier */
for( j=0; sym[i].id[j++] != '\0'; ) ...
etc.
```

Thus to print a list of all identifiers that haven't been used, together with their line number,

```
for( i=0; i<nsym; i++ )
    if( sym[i].usage == 0 )
        printf("%d\t%s\n", sym[i].line, sym[i].id);
```

Suppose we now want to write a function `lookup(name)` which will tell us if `name` already exists in `sym`, by giving its index, or that it doesn't, by returning a `-1`. We can't pass a structure to a function directly — we have to either define it externally, or pass a pointer to it. Let's try the first way first.

```
int    nsym 0; /* current length of symbol table */
struct {
    char  id[10];
    int   line;
    char  type;
    int   usage;
} sym[100]; /* symbol table */

main() {
    ...
    if( (index = lookup(newname)) >= 0 )
        sym[index].usage++; /* already there ... */
    else
        install(newname, newline, newtype);
    ...
}

lookup(s)
char *s; {
    int i;
    extern struct {
        char  id[10];
        int   line;
        char  type;
        int   usage;
    } sym[ ];
    for( i=0; i<nsym; i++ )
        if( compar(s, sym[i].id) > 0 )
            return(i);
    return(-1);
}

compar(s1,s2) /* return 1 if s1==s2, 0 otherwise */
char *s1, *s2; {
    while( *s1++ == *s2 )
        if( *s2++ == '\0' )
            return(1);
}
```

```

        return(0);
    }

```

The declaration of the structure in `lookup` isn't needed if the external definition precedes its use in the same source file, as we shall see in a moment.

Now what if we want to use pointers?

```

struct symtag {
    char  id[10];
    int   line;
    char  type;
    int   usage;
} sym[100], *psym;

psym = &sym[0];    /* or psym = sym; */

```

This makes `psym` a pointer to our kind of structure (the symbol table), then initializes it to point to the first element of `sym`.

Notice that we added something after the word `struct`: a "tag" called `symtag`. This puts a name on our structure definition so we can refer to it later without repeating the definition. It's not necessary but useful. In fact we could have said

```

struct symtag {
    ... structure definition
};

```

which wouldn't have assigned any storage at all, and then said

```

struct symtag sym[100];
struct symtag *psym;

```

which would define the array and the pointer. This could be condensed further, to

```

struct symtag sym[100], *psym;

```

The way we actually refer to an member of a structure by a pointer is like this:

```
ptr -> structure-member
```

The symbol `->` means we're pointing at a member of a structure; `->` is only used in that context. `ptr` is a pointer to the (base of) a structure that contains the structure member. The expression `ptr->structure-member` refers to the indicated member of the pointed-to structure. Thus we have constructions like:

```

psym->type = 1;
psym->id[0] = 'a';

```

and so on.

For more complicated pointer expressions, it's wise to use parentheses to make it clear who goes with what. For example,

```

struct { int x, *y; } *p;
p->x++      increments x
++p->x     so does this!
(++p)->x   increments p before getting x
*p->y++    uses y as a pointer, then increments it
*(p->y)++  so does this
*(p++)->y  uses y as a pointer, then increments p

```

The way to remember these is that `->`, `.` (dot), `()` and `[]` bind very tightly. An expression in-

volving one of these is treated as a unit. $p \rightarrow x$, $a[i]$, $y.x$ and $f(b)$ are names exactly as abc is.

If p is a pointer to a structure, any arithmetic on p takes into account the actual size of the structure. For instance, $p++$ increments p by the correct amount to get the next element of the array of structures. But don't assume that the size of a structure is the sum of the sizes of its members — because of alignments of different sized objects, there may be "holes" in a structure.

Enough theory. Here is the lookup example, this time with pointers.

```

struct symtag {
    char   id[10];
    int    line;
    char   type;
    int    usage;
} sym[100];

main( ) {
    struct symtag *lookup( );
    struct symtag *psym;
    ...
    if( (psym = lookup(newname)) ) /* non-zero pointer */
        psym -> usage++;         /* means already there */
    else
        install(newname, newline, newtype);
    ...
}

struct symtag *lookup(s)
char *s; {
    struct symtag *p;
    for( p=sym; p < &sym[nsym]; p++ )
        if( compar(s, p->id) > 0)
            return(p);
    return(0);
}

```

The function `compar` doesn't change: $p \rightarrow id$ refers to a string.

In `main` we test the pointer returned by `lookup` against zero, relying on the fact that a pointer is by definition never zero when it really points at something. The other pointer manipulations are trivial.

The only complexity is the set of lines like

```
struct symtag *lookup( );
```

This brings us to an area that we will treat only hurriedly — the question of function types. So far, all of our functions have returned integers (or characters, which are much the same). What do we do when the function returns something else, like a pointer to a structure? The rule is that any function that doesn't return an `int` has to say explicitly what it does return. The type information goes before the function name (which can make the name hard to see). Examples:

```

char f(a)
    int a; {
    ...
}

```

```
int *g() { ... }
```

```
struct symtag *lookup(s) char *s; { ... }
```

The function `f` returns a character, `g` returns a pointer to an integer, and `lookup` returns a pointer to a structure that looks like `symtag`. And if we're going to use one of these functions, we have to make a declaration where we use it, as we did in `main` above.

Notice the parallelism between the declarations

```
struct symtag *lookup( );
struct symtag *psym;
```

In effect, this says that `lookup()` and `psym` are both used the same way — as a pointer to a structure — even though one is a variable and the other is a function.

21. Initialization of Variables

An external variable may be initialized at compile time by following its name with an initializing value when it is defined. The initializing value has to be something whose value is known at compile time, like a constant.

```
int    x    0;    /* "0" could be any constant */
int    a    'a';
char   flag  0177;
int    *p    &y[1]; /* p now points to y[1] */
```

An external array can be initialized by following its name with a list of initializations enclosed in braces:

```
int    x[4]  {0,1,2,3};    /* makes x[i] = i */
int    y[ ]  {0,1,2,3};    /* makes y big enough for 4 values */
char   *msg  "syntax error\n"; /* braces unnecessary here */
char *keyword[ ]{
    "if",
    "else",
    "for",
    "while",
    "break",
    "continue",
    0
};
```

This last one is very useful — it makes `keyword` an array of pointers to character strings, with a zero at the end so we can identify the last element easily. A simple `lookup` routine could scan this until it either finds a match or encounters a zero keyword pointer:

```
lookup(str)    /* search for str in keyword[ ] */
char *str; {
    int i,j,r;
    for( i=0; keyword[i] != 0; i++) {
        for( j=0; (r=keyword[i][j]) == str[j] && r != '\0'; j++ );
        if( r == str[j] )
            return(i);
    }
    return(-1);
}
```

Sorry — neither local variables nor structures can be initialized.

22. Scope Rules: Who Knows About What

A complete C program need not be compiled all at once; the source text of the program may be kept in several files, and previously compiled routines may be loaded from libraries. How do we arrange that data gets passed from one routine to another? We have already seen how to use function arguments and values, so let us talk about external data. Warning: the words *declaration* and *definition* are used precisely in this section; don't treat them as the same thing.

A major shortcut exists for making **extern** declarations. If the definition of a variable appears *before* its use in some function, no **extern** declaration is needed within the function. Thus, if a file contains

```
f1() { ... }
int foo;
f2() { ... foo = 1; ... }
f3() { ... if ( foo ) ... }
```

no declaration of **foo** is needed in either **f2** or **f3**, because the external definition of **foo** appears before them. But if **f1** wants to use **foo**, it has to contain the declaration

```
f1() {
    extern int foo;
    ...
}
```

This is true also of any function that exists on another file — if it wants **foo** it has to use an **extern** declaration for it. (If somewhere there is an **extern** declaration for something, there must also eventually be an external definition of it, or you'll get an "undefined symbol" message.)

There are some hidden pitfalls in external declarations and definitions if you use multiple source files. To avoid them, first, define and initialize each external variable only once in the entire set of files:

```
int    foo    0;
```

You can get away with multiple external definitions on UNIX, but not on GCOS, so don't ask for trouble. Multiple initializations are illegal everywhere. Second, at the beginning of any file that contains functions needing a variable whose definition is in some other file, put in an **extern** declaration, outside of any function:

```
extern int    foo;
f1() { ... }
etc.
```

The **#include** compiler control line, to be discussed shortly, lets you make a single copy of the external declarations for a program and then stick them into each of the source files making up the program.

23. #define, #include

C provides a very limited macro facility. You can say

```
#define    name    something
```

and thereafter anywhere "name" appears as a token, "something" will be substituted. This is

particularly useful in parametering the sizes of arrays:

```
#define    ARRAYSIZE 100
int    arr[ARRAYSIZE];
...
while( i++ < ARRAYSIZE )...
```

(now we can alter the entire program by changing only the **define**) or in setting up mysterious constants:

```
#define    SET        01
#define    INTERRUPT  02    /* interrupt bit */
#define    ENABLED    04
...
if( x & (SET | INTERRUPT | ENABLED) ) ...
```

Now we have meaningful words instead of mysterious constants. (The mysterious operators '&' (AND) and '|' (OR) will be covered in the next section.) It's an excellent practice to write programs without any literal constants except in **#define** statements.

There are several warnings about **#define**. First, there's no semicolon at the end of a **#define**; all the text from the name to the end of the line (except for comments) is taken to be the "something". When it's put into the text, blanks are placed around it. Good style typically makes the name in the **#define** upper case — this makes parameters more visible. Definitions affect things only after they occur, and only within the file in which they occur. Defines can't be nested. Last, if there is a **#define** in a file, then the first character of the file *must* be a '#', to signal the preprocessor that definitions exist.

The other control word known to C is **#include**. To include one file in your source at compilation time, say

```
#include "filename"
```

This is useful for putting a lot of heavily used data definitions and **#define** statements at the beginning of a file to be compiled. As with **#define**, the first line of a file containing a **#include** has to begin with a '#'. And **#include** can't be nested — an included file can't contain another **#include**.

24. Bit Operators

C has several operators for logical bit-operations. For example,

```
x = x & 0177;
```

forms the bit-wise AND of *x* and 0177, effectively retaining only the last seven bits of *x*. Other operators are

	inclusive OR
^	(circumflex) exclusive OR
~	(tilde) 1's complement
!	logical NOT
<<	left shift (as in $x \ll 2$)
>>	right shift (arithmetic on PDP-11; logical on H6070, IBM360)

25. Assignment Operators

An unusual feature of C is that the normal binary operators like '+', '-', etc. can be combined with the assignment operator '=' to form new assignment operators. For example,

```
x -= 10;
```

uses the assignment operator '=-' to decrement x by 10, and

```
x = & 0177
```

forms the AND of x and 0177. This convention is a useful notational shortcut, particularly if x is a complicated expression. The classic example is summing an array:

```
for( sum=i=0; i<n; i++ )
    sum =+ array[i];
```

But the spaces around the operator are critical! For instance,

```
x = - 10;
```

sets x to -10, while

```
x = - 10;
```

subtracts 10 from x. When no space is present,

```
x = - 10;
```

also decreases x by 10. This is quite contrary to the experience of most programmers. In particular, watch out for things like

```
c = *s + +;
y = &x[0];
```

both of which are almost certainly not what you wanted. Newer versions of various compilers are courteous enough to warn you about the ambiguity.

Because all other operators in an expression are evaluated before the assignment operator, the order of evaluation should be watched carefully:

```
x = x << y | z;
```

means "shift x left y places, then OR with z, and store in x." But

```
x = << y | z;
```

means "shift x left by y|z places", which is rather different.

26. Floating Point

We've skipped over floating point so far, and the treatment here will be hasty. C has single and double precision numbers (where the precision depends on the machine at hand). For example,

```
double sum;
float avg, y[10];
sum = 0.0;
for( i=0; i<n; i++ )
    sum =+ y[i];
avg = sum/n;
```

forms the sum and average of the array y.

All floating arithmetic is done in double precision. Mixed mode arithmetic is legal; if an arithmetic operator in an expression has both operands int or char, the arithmetic done is integer, but if one operand is int or char and the other is float or double, both operands are con-

verted to **double**. Thus if *i* and *j* are **int** and *x* is **float**,

```
(x+i)/j    converts i and j to float
x + i/j    does i/j integer, then converts
```

Type conversion may be made by assignment; for instance,

```
int m, n;
float x, y;
m = x;
y = n;
```

converts *x* to integer (truncating toward zero), and *n* to floating point.

Floating constants are just like those in Fortran or PL/I, except that the exponent letter is 'e' instead of 'E'. Thus

```
pi = 3.14159;
large = 1.23456789e10;
```

printf will format floating point numbers: "%w.df" in the format string will print the corresponding variable in a field *w* digits wide, with *d* decimal places. An **e** instead of an **f** will produce exponential notation.

27. Horrors! goto's and labels

C has a **goto** statement and labels, so you can branch about the way you used to. But most of the time **goto**'s aren't needed. (How many have we used up to this point?) The code can almost always be more clearly expressed by **for/while**, **if/else**, and compound statements.

One use of **goto**'s with some legitimacy is in a program which contains a long loop, where a **while(1)** would be too extended. Then you might write

```
mainloop:
...
goto mainloop;
```

Another use is to implement a **break** out of more than one level of **for** or **while**. **goto**'s can only branch to labels within the same function.

28. Acknowledgements

I am indebted to a veritable host of readers who made valuable criticisms on several drafts of this tutorial. They ranged in experience from complete beginners through several implementors of C compilers to the C language designer himself. Needless to say, this is a wide enough spectrum of opinion that no one is satisfied (including me); comments and suggestions are still welcome, so that some future version might be improved.


References

C is an extension of B, which was designed by D. M. Ritchie and K. L. Thompson [4]. The C language design and UNIX implementation are the work of D. M. Ritchie. The GCOS version was begun by A. Snyder and B. A. Barres, and completed by S. C. Johnson and M. E. Lesk. The IBM version is primarily due to T. G. Peterson, with the assistance of M. E. Lesk.

- [1] D. M. Ritchie, *C Reference Manual*. Bell Labs, Jan. 1974.
- [2] M. E. Lesk & B. A. Barres, *The GCOS C Library*. Bell Labs, Jan. 1974.
- [3] D. M. Ritchie & K. Thompson, *UNIX Programmer's Manual*. 5th Edition, Bell Labs, 1974.
- [4] S. C. Johnson & B. W. Kernighan, *The Programming Language B*. Computer Science Technical Report 8, Bell Labs, 1972.







A New Input-Output Package

D. M. Ritchie

Bell Laboratories
Murray Hill, New Jersey 07974

A new package of IO routines is available. It was designed with the following goals in mind.

1. It should be similar in spirit to the earlier Portable Library, and, to the extent possible, be compatible with it. At the same time a few dubious design choices in the Portable Library will be corrected.
2. It must be as efficient as possible, both in time and in space, so that there will be no hesitation in using it no matter how critical the application.
3. It must be simple to use, and also free of the magic numbers and mysterious calls the use of which mars the understandability and portability of many programs using older packages.
4. The interface provided should be applicable on all machines, whether or not the programs which implement it are directly portable to other systems, or to machines other than the PDP11 running a version of Unix.

It is intended that this package replace the Portable Library. Although it is not directly compatible, as discussed below, it is sufficiently similar that modifying programs to use it should be a simple exercise.

The most crucial difference between this package and the Portable Library is that the current offering names streams in terms of pointers rather than by the integers known as 'file descriptors.' Thus, for example, the routine which opens a named file returns a pointer to a certain structure rather than a number; the routine which reads an open file takes as an argument the pointer returned from the open call.

General Usage

Each program using the library must have the line

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

which defines certain macros and variables. The library containing the routines is '/usr/lib/libS.a,' so the command to compile is

```
cc ... -lS
```

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

stdin	The name of the standard input file
stdout	The name of the standard output file
stderr	The name of the standard error file
EOF	is actually -1, and is the value returned by the read routines on end-of-file or error.

NULL is a notation for the null pointer, returned by pointer-valued functions to indicate an error

FILE expands to 'struct _iob' and is a useful shorthand when declaring pointers to streams.

BUFSIZ is a number (viz. 512) of the size suitable for an IO buffer supplied by the user. See *setbuf*, below.

getc, *getchar*, *putc*, *putchar*, *feof*, *ferror*, *fileno* are defined as macros. Their actions are described below; they are mentioned here to point out that it is not possible to redeclare them and that they are not actually functions; thus, for example, they may not have breakpoints set on them.

The routines in this package, like the Portable Library, offer the convenience of automatic buffer allocation and output flushing where appropriate. Absent, however, is the facility of changing the default input and output streams by assigning to 'cin' and 'cout.' The names 'stdin,' 'stdout,' and 'stderr' are in effect constants and may not be assigned to.

Calls

The routines in the library are in nearly one-to-one correspondence with those in the Portable Library. In several cases the name has been changed. This is an attempt to reduce confusion.

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

Fopen opens the file and, if needed, allocates a buffer for it. *Filename* is a character string specifying the name. *Type* is a character string (not a single character). It may be "r", "w", or "a" to indicate intent to read, write, or append. The value returned is a file pointer. If it is NULL the attempt to open failed.

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

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

*int getc(ioptr) FILE *ioptr*

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

*int fgetc(ioptr) FILE *ioptr*

acts like *getc* but is a genuine function, not a macro.

*putc(c, ioptr) FILE *ioptr*

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

*fputc(c, ioptr) FILE *ioptr*

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

*fclose(ioptr) FILE *ioptr*

The file corresponding to *ioptr* is closed after any buffers are emptied. A buffer allocated by the IO system is freed. *Fclose* is automatic on normal termination of the program.

*fflush(ioptr) FILE *ioptr*

Any buffered information on the (output) stream named by *ioptr* is written out. Output files are normally buffered if and only if they are not directed to the terminal, but *stderr* is unbuffered unless *setbuf* is used.

exit(errcode)

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

*feof(ioptr) FILE *ioptr*

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

*ferror(ioptr) FILE *ioptr*

returns non-zero when an error has occurred while reading or writing the named stream. The error indication lasts until the file has been closed.

getchar()

is identical to *getc(stdin)*.

putchar(c)

is identical to *putc(c, stdout)*.

*char *gets(s) char *s*

reads characters up to a new-line from the standard input. The new-line character is replaced by a null character. It is the user's responsibility to make sure that the character array *s* is large enough. *Gets* returns its argument, or NULL if end-of-file or error occurred. Note that this routine is not compatible with *fgets*; it is included for downward compatibility.

*char *fgets(s, n, ioptr) char *s; FILE *ioptr*

reads up to *n* characters from the stream *ioptr* into the character pointer *s*. The read terminates with a new-line character. The new-line character is placed in the buffer followed by a null character. The first argument, or NULL if error or end-of-file occurred, is returned.

*puts(s) char *s*

writes the null-terminated string (character array) *s* on the standard output. A new-line is appended. No value is returned. Note that this routine is not compatible with *fputs*; it is included for downward compatibility.

**fputs(s, ioptr) char *s; FILE *ioptr*

writes the null-terminated string (character array) *s* on the stream *ioptr*. No new-line is appended. No value is returned.

*ungetc(c, ioptr) FILE *ioptr*

The argument character *c* is pushed back on the input stream named by *ioptr*. Only one character may be pushed back.

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

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

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

Printf writes on the standard output. *Fprintf* writes on the named output stream. *Sprintf* puts characters in the character array (string) named by *s*. The specifications are as described in section *printf* (III) of the Unix Programmer's Manual. There is a new conversion: *%m.ng* converts a double argument in the style of *e* or *f* as most appropriate.

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

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

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

Scanf reads from the standard input. *Fscanf* reads from the named input stream. *Sscanf* reads from the character string supplied as *s*. The specifications are identical to those of the Portable Library. *Scanf* reads characters, interprets them according to a format, and stores the results in its arguments. It expects as arguments a control string *format*, described below, and a set of

arguments, *each of which must be a pointer*, indicating where the converted input should be stored.

The control string usually contains conversion specifications, which are used to direct interpretation of input sequences. The control string may contain:

1. **Blanks**, tabs or newlines, which are ignored.
2. Ordinary characters (not %) which are expected to match the next non-space character of the input stream (where space characters are defined as blank, tab or newline).
3. Conversion specifications, consisting of the character %, an optional assignment suppressing character, an optional numerical maximum field width, and a conversion character.

A conversion specification is used to direct the conversion of the next input field; the result is placed in the variable pointed to by the corresponding argument, unless assignment suppression was indicated by the character. An input field is defined as a string of non-space characters; it extends either to the next space character or until the field width, if specified, is exhausted.

The conversion character indicates the interpretation of the input field; the corresponding pointer argument must usually be of a restricted type. The following conversion characters are legal:

- % indicates that a single % character is expected in the input stream at this point; no assignment is done.
- d indicates that a decimal integer is expected in the input stream; the corresponding argument should be an integer pointer.
- o indicates that an octal integer is expected in the input stream; the corresponding argument should be a integer pointer.
- x indicates that a hexadecimal integer is expected in the input stream; the corresponding argument should be an integer pointer.
- s indicates that a character string is expected; the corresponding argument should be a character pointer pointing to an array of characters large enough to accept the string and a terminating '\0', which will be added. The input field is terminated by a space character or a newline.
- c indicates that a character is expected; the corresponding argument should be a character pointer; the next input character is placed at the indicated spot. The normal skip over space characters is suppressed in this case; to read the next non-space character, try %Is. If a field width is given, the corresponding argument should refer to a character array, and the indicated number of characters is read.
- e (or f) indicates that a floating point number is expected in the input stream; the next field is converted accordingly and stored through the corresponding argument, which should be a pointer to a *float*. The input format for floating point numbers is an optionally signed string of digits possibly containing a decimal point, followed by an optional exponent field beginning with an E or e followed by an optionally signed integer.
- [indicates a string not to be delimited by space characters. The left bracket is followed by a set of characters and a right bracket; the characters between the brackets define a set of characters making up the string. If the first character is not circumflex (^), the input field is all characters until the first character not in the set between the brackets; if the first character after the left bracket is ^, the input field is all characters until the first character which is in the remaining set of characters between the brackets. The corresponding argument must point to a character array.

The conversion characters *d*, *o* and *x* may be capitalized or preceded by *l* to indicate that a pointer to *long* rather than *int* is expected. Similarly, the conversion characters *e* or *f* may be capitalized or preceded by *l* to indicate that a pointer to *double* rather than *float* is in the argument list. The character *h* will function similarly in the future to indicate *short* data items.

For example, the call

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

with the input line

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

will assign to *i* the value 25, *x* the value 5.432, and *name* will contain "thompson\0". Or,

```
int i; float x; char name[50];
scanf("%2d%f%d%[1234567890]", &i, &x, name);
```

with input

```
56789 0123 56a72
```

will assign 56 to *i*, 789.0 to *x*, skip "0123", and place the string "56\0" in *name*. The next call to *getchar* will return 'a'.

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

*fread(ptr, sizeof(*ptr), nitems, ioptr) FILE *ioptr*

reads *nitems* of data beginning at *ptr* from file *ioptr*. It behaves identically to the Portable Library's *cread*. No advance notification that binary IO is being done is required; when, for portability reasons, it becomes required, it will be done by adding an additional character to the mode-string on the *fopen* call.

*fwrite(ptr, sizeof(*ptr), nitems, ioptr) FILE *ioptr*

Like *fread*, but in the other direction.

*rewind(ioptr) FILE *ioptr*

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

*system(string) char *string*

The *string* is executed by the shell as if typed at the terminal.

*getw(ioptr) FILE *ioptr*

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

*putw(w, ioptr) FILE *ioptr*

writes the integer *w* on the named output stream.

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

Setbuf may be used after a stream has been opened but before IO has started. If *buf* is NULL, the stream will be unbuffered. Otherwise the buffer supplied will be used. It is a character array of sufficient size:

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

*fileno(ioptr) FILE *ioptr*

returns the integer file descriptor associated with the file.

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

The location of the next byte in the stream named by *ioptr* is adjusted. *Offset* is a long integer. If *ptrname* is 0, the offset is measured from the beginning of the file; if *ptrname* is 1, the offset

is measured from the current read or write pointer; if *ptrname* is 2, the offset is measured from the end of the file. The routine accounts properly for any buffering. (When this routine is used on non-Unix systems, the offset must be a value returned from *ftell* and the *ptrname* must be 0).

*long ftell(ioptr) FILE *ioptr*

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

*getpw(uid, buf) char *buf*

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

*char *calloc(num, size)*

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

*cfree(ptr) char *ptr*

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

The following are macros defined by *stdio.h*.

isalpha(c)

returns non-zero if the argument is alphabetic.

isupper(c)

returns non-zero if the argument is upper-case alphabetic.

islower(c)

returns non-zero if the argument is lower-case alphabetic.

isdigit(c)

returns non-zero if the argument is a digit.

isspace(c)

returns non-zero if the argument is a spacing character: tab, new-line, carriage return, vertical tab, form feed, space.


toupper(c)

returns the upper-case character corresponding to the lower-case letter *c*.

tolower(c)

returns the lower-case character corresponding to the upper-case letter *c*.





A General-Purpose Subroutine Library for PWB/UNIX

Alan L. Glasser

Bell Laboratories
Piscataway, New Jersey 08854

ABSTRACT

A new general-purpose subroutine library has been written for PWB/UNIX. It complements the functions provided by D. M. Ritchie's *A New Input-Output Package*. This library has been used in the implementation of Release 4 of the Programmer's Workbench Source Code Control System (SCCS/PWB), and many small UNIX C programs. It is efficient in time and space, as well as being easy to use. This document is a user's guide to this library.

1. "Include" FILES

The directory `"/usr/include"` contains public *include* files. Users of the subroutine library should be familiar with the contents of these files. The following are available:

archive.h	Defines the "magic" number of an archive file and declares a structure for the header of an archive file.
ctype.h	Defines macros for testing whether a character is alphabetic, upper-case, lower-case, a digit, a "space," and for converting upper-case characters to lower-case and vice-versa.
dir.h	Declares a structure for a directory entry.
errnos.h	Defines system call error numbers (see <i>INTRO(II)</i>). These were copied from the UNIX system source.
fatal.h	Defines certain macros, constants, and variables used by the general-purpose error and signal handling subroutines (see below).
macros.h	Defines some general-purpose macros.
misc.h	Declares some unnamed structures for accessing pieces of a variable (e.g., the low byte of an integer). These were copied from the UNIX system source.
stat.h	Declares an inode structure to be used with <i>stat(II)</i> and <i>fstat(II)</i> , and defines constants for the various mode bits of an inode.
stdio.h	Used by <i>A New Input-Output Package</i> .
time.h	Declares a structure to be used with <i>localtime(III)</i> .
system.h	Defines certain system constants; in particular, signal numbers. These were copied from the UNIX system source.

2. SUBROUTINES

There are four sets of subroutines available. Three of these sets are kept in one library (`"/usr/lib/libpw.a"`, accessible as `-lpw`), and the fourth set is kept in a second library (`"/usr/lib/libwrt.a"`, accessible as `-lwrt`) for reasons that will be explained below; `-lwrt` should normally be last on the `cc` argument list.

The first library (`-lpw`) contains the *string*, *error*, and *sys* sets. The second library (`-lwrt`) contains only the *write* set (see below).

The *string* set can be used in conjunction with any other subroutines. The *string* set is independent of the other sets and of other subroutine libraries; no subroutine in the *string* set calls any subroutine not in the *string* set.

The *error* set provides general-purpose error and signal handling routines.

The *sys* set provides interfaces to most of the commonly used system calls. These interfaces will call *fatal()* (see the *error* set) if an error condition is detected in the interface.

The *write* set contains an interface to the *write(II)* system call; this interface handles errors and calls *fatal()*, if necessary. If the *write* set is used, then all calls to the subroutine *write()* will be directed to this *write* routine (remember, it will call *fatal()* if an error is detected). It is for this reason that the *write* set is in a separate library; if all routines were in the same library, then users would be unwittingly using the *write(II)* routine of the *write* set.

The subroutines presented here do not call *nargs(III)*, as the *nargs(III)* subroutine does not work if a program is loaded with separate data and instruction spaces.

2.1 String Set

*char *alloca(nbytes)*

Allocates *nbytes* bytes of automatic memory. Automatic memory is freed upon return from the calling function (space is allocated from the stack). There is no way to *explicitly* free a piece of memory gotten from *alloca()*. *Alloca()* returns the address of the allocated area; a memory fault is generated when there isn't enough memory.

N. B.: Use of *alloca()* as an argument to another subroutine will *not* work correctly because arguments are pushed onto the stack and *alloca()* takes memory from the stack. It is necessary to first set a temporary variable equal to the value returned by *alloca()*, and then pass that variable to the other subroutine.

any(c, str)

If character *c* is equal to any character in the string *str*, returns 1; else returns 0.

anystr(str1, str2)

If any character of string *str1* is equal to any character in string *str2*, returns the offset (in *str1*) of the first such match; else returns -1.

balbrk(str, open, clos, end)

Finds the offset, in string *str*, of the first of the characters in the string *end* occurring *outside* of a *balanced string*. A balanced string contains matched occurrences of any character in the string *open* and the corresponding character in the string *clos*. Balanced strings may be nested. In addition to the characters in *end*, the null character is implicitly an *end* character. Unmatched members of *open* or *close* result in an error return (a value of -1 is returned).

Example 1:

```
s = "a[bc=2]=3";
o = "([";
c = ")]";
e = "=";
```

balbrk(s, o, c, e) returns 7.

Example 2:

```
s = "a[bc=2]=3";
```

with *o*, *c*, and *e* as in Example 1, *balbrk(s, o, c, e)* returns -1.

*char *cat(dest, source₁, source₂, source₃, ..., source_n, 0)*

Concatenates strings. First, string *source₁* is copied to string *dest*. Then subsequent *source_k* strings are concatenated (by copying) onto the end of *dest*. The space for *dest* must be allocated by the caller (i.e., *dest* is taken to be the address of an area of memory large enough to hold the result). The address of the result (i.e., *dest*) is returned.

dname(pathname)

Returns a pointer to the name of the directory that contains the file pointed to by *pathname*. *Dname()* is the complement of *sname()* (see below). If *pathname* is a simple name (e.g., "file"), a pointer to "." is returned. If *pathname* is "/unix", a pointer to "/" is returned. If *pathname* is "/bin/who", a pointer to "/bin" is returned, etc. The string pointed to by *pathname* is modified by *dname()*; *pathname* is returned.

equal(str1, str2)

If string *str1* is equal to string *str2*, returns 1; else returns 0.

imatch(prefix, str)

Initial match. If string *prefix* is a prefix of string *str*, returns 1; else returns 0.

index(str1, str2)

If string *str2* is a substring of string *str1*, returns the offset of the first occurrence of *str2* in *str1*, else returns -1.

move(a, b, n)

Copies the first *n* characters from string *a* to string *b*.

atoi(str)

Converts an ASCII string to a integer. The string *str* is taken to be a string of decimal digits; the numeric value represented by *str* is returned. Converts positive numbers only. Returns -1 if a non-numeric character is encountered.

long atoi(str)

Converts an ASCII string to a long integer. The string *str* is taken to be a string of decimal digits; the numeric value represented by *str* is returned. Converts positive numbers only. Returns -1 if a non-numeric character is encountered.

*char *repeat(result, str, repfac)*

The string *str* is first copied to the string *result*. Then *str* is copied *repfac*-1 times onto the end of *result*. As with *cat()* (see above), allocation of space for *result* is the caller's responsibility. *Result* is returned.

*char *satoi(str, ip)*

Satoi() is similar to *atoi* (see above), except that the integer value is stored through the integer pointer *ip*, and a pointer to the first non-numeric character encountered is returned.

size(str)

Returns the number of bytes of memory used by string *str*, including the null byte, so that *size(str)* is equal to *length(str)*+1.

*char *sname(str)*

Sname() returns a pointer to the "simple" name of path name *str*; i.e., a pointer to the first character after the last "/" in *str*. If *str* does not contain a "/", a pointer to the original string is returned.

*char *strend(str)*

Strend() returns a pointer to the end (null byte) of the string *str*.

substr(str, result, origin, len)

Copies at most *len* characters from the string *str* starting at *str[origin]* to the string pointed to by *result*. Sufficient space must exist for that string; *result* is returned. There is *no checking* for the reasonableness of the arguments. The copying of *str* to *result* stops if either the specified number (i.e., *len*, which is taken as an unsigned integer) characters have been copied, or if the end of *str* (i.e., a null byte) is found. A large value of *len* (e.g., -1) will usually cause all of *str* to be copied.

*char *trnslat(str, old, new, result)*

Copies string *str* to string *result* replacing any character found in string *old* with the corresponding character from string *new*; *result* is returned.

verify(str1, str2)

If string *str1* contains any characters not in string *str2*, returns the offset of the first such character in *str1*; else returns -1.

*char *zero(ptr, cnt)*

Sets to zero the area of memory *cnt* bytes long, starting at address *ptr*; *ptr* is returned.

*char *zeropad(str)*

Replace initial blanks with "0" characters in string *str*; *str* is returned.

2.2 Error Set

The *error* set of subroutines consists of a general-purpose error handling routine called *fatal()*, and general-purpose signal-setting and signal-catching routines called *setsig()* and *setsig1()*, respectively. There are also two additional routines called *clean_up()* and *userexit()*, which may be called by *fatal()* or *setsig1()*. Default versions of these two additional routines are supplied in the library. Users may define their own *clean_up()* and *userexit()* routines.

The public *include* file "fatal.h" contains definitions needed to use *fatal()*. It contains the following:

```
extern int    Fflags;
extern char  *Ffile;
extern int    Fvalue;
extern int    (*Ffunc)();
extern int    Fjmp[3];

# define FTLMSG    0100000
# define FTLCLN   040000
# define FTLFUNC  020000
# define FTLACT   077
# define FTLJMP   02
# define FTLEXIT  01
# define FTLRET   0

# define FSAVE(val) SAVE(Fflags,old_Fflags); Fflags = val;
# define FRSTR()    RSTR(Fflags,old_Fflags);
```

fatal(msg)

A general-purpose error handler. Typically, low-level subroutines that detect error conditions (an open or create routine, for example) return as a value a call of *fatal()* with an appropriate message string. For example:

```
return(fatal("can't do it"));
```

Higher-level routines control the execution of *fatal()* via the global word *Fflags*. The macros *FSAVE()* and *FRSTR()* in "fatal.h" can be used by higher-level subroutines to save and restore the *Fflags* word.

The argument to *fatal()* is a pointer to a message string. The action of *fatal()* is driven completely from the *Fflags* global integer, which is interpreted as explained below.

The *FTLMSG* bit controls the writing of the message on file descriptor 2. The message is preceded by the string "ERROR: ", unless the global character pointer *Ffile* is non-zero, in which case the message is preceded by a string equivalent to:

```
s = sprintf(space, "ERROR [%s]: ", Ffile);
```

A new-line character is written after the user-supplied message.

If the *FTLCLN* bit is on, *clean_up()* is called with an argument of 0 (see below).

If the *FTLFUNC* bit is on, the function pointed to by the global function pointer *Ffunc* is called with the user-supplied message pointer as an argument. This feature can be used to log these messages.

The *FTLACT* bits determine how *fatal()* should return. If the *FTLJMP* bit is one, *longjmp(Fjmp)* (see *setjmp(III)*) is called. If the *FTLEXIT* bit is one the value of *userexit(1)* is passed as an argument to

exit(II) (see below). If none of the FTLACT bits is on (the default value for *Fflags* is 0), the global word *Fvalue* (initialized to -1) is returned.

If all *fatal()* globals have their default values, *fatal()* simply returns -1.

setSIG()

General-purpose signal-setting routine. All signals not already ignored or caught are made to be caught by the signal catching routine *setSIG()*.

setSIG()

General-purpose signal catcher and termination routine. If a signal other than hangup, interrupt, or quit is caught, a "user-oriented" *help(I)* message is printed on file descriptor 2. If hangup, interrupt, or quit is caught, subsequent occurrences of that signal will be ignored. Termination is similar to the FTLCLN and FTLEXIT options of *fatal()*, in that *clean_up(sig)* (where *sig* is the signal number) and *exit(userexit(I))* are called.

If the file "dump.core" exists in the current directory, the IOT signal is set to 0 and *abort(III)* is called to produce a core dump (after calling *clean_up()*, but before calling *userexit()*).

clean_up()

A default *clean_up()* routine is provided to resolve external references. It simply returns. User-supplied *clean_up()* routines are often used for removing temporary files, etc.

userexit(code)

A default *userexit()* routine is provided to resolve external references. It returns the value of *code*. User-supplied *userexit()* routines are often used for logging usage statistics.

2.3 Sys Set

The *sys* set of subroutines provides interfaces to system calls that process error conditions and call *fatal()*. In addition, a few functions which are not available elsewhere are provided.

curdir(path)

Places the complete pathname of the current directory in string *path*. Returns 0 on success, non-zero on failure. On successful return, the current directory is the same as it was on entry; on failure return, the current directory is not known.

jdlopen(fd, mode)

This subroutine provides a file-descriptor interface to the routines in *A New Input-Output Package*, and is required when one wants to use the routines in *A New Input-Output Package* with pipes. The first argument is a file descriptor (from *open(II)*, *creat(II)*, or *pipe(II)*), the second is the read/write mode (0/1, respectively). A *file pointer* (see *A New Input-Output Package*) is returned on success, and NULL on failure (typically, because there are no file structures available).

giveup(dump)

This routine does the following:

- Change directory to "/" if argument is 0.
- Set IOT signal to system default (0).
- Call *abort(III)*.

Thus, if *giveup()* is called with a 0 argument, and the file "/core" is not writable (or if the file "/core" doesn't exist, and the directory "/" is not writable), no core dump will be produced.

lockit(lockfile, count, pid)

A process semaphore implemented with files; typically, used to establish exclusive use of a resource (usually a file). The file's name is *lockfile*. *Lockit()* tries *count* times to create *lockfile* mode 444. It sleeps 10 seconds between tries. If *lockfile* is created, the number *pid* (typically, the process ID of the current process) is written (in binary; i.e., as two bytes) into *lockfile*, and 0 is returned. If *lockfile* exists and hasn't been modified within the last 60 seconds, and if it either is empty or if its first two bytes, interpreted as a binary number, are *not* the process ID of any existing process, *lockfile* is removed and *lockit()* tries again to make *lockfile*. After *count* tries, or if the reason for the creation of *lockfile* failing is something other than EACCES (see *INTRO(II)*), *lockit()* returns -1. See also *unlockit()*, below.

rename(oldname, newname)

Renames *oldname* to be *newname*; it can be thought of as:

```
mv oldname newname
```

It calls *xlink()* and *xunlink()* (see below).

unlockit(lockfile, pid)

Unlockit() is meant to be used to remove a *lockfile* created by *lockit()*. It verifies that the *pid* specified is contained in the first two bytes of the named *lockfile*, and then removes it. If the *pids* match, and the file is successfully removed, *unlockit()* returns 0; otherwise, -1 is returned.

userdir(uid)

Returns user's login directory name. The argument must be an integer user ID. There is an assumption that the directory field is the fifth field of a password file entry (i.e., there is no "group id" in the password file). Returns a pointer to the login directory on success, 0 on failure. It remembers its argument and the returned login directory name for subsequent calls to speed itself up. Users of PWB/UNIX systems should use *logdir()* (see *loginfo(II)*).

username(uid)

Returns user's login name. The argument must be an integer user ID. Returns a pointer to the login name on success, a pointer to the string representation of the user ID on failure. There is an assumption that the login name field is the first field of a password file entry. It remembers its argument and the returned login name for subsequent calls to speed itself up. Users of PWB/UNIX systems should use *logname()* (see *loginfo(II)*).

xalloc(size), xfree(ptr), xfreeall()

Xalloc() and *xfree()* are used in the same way as *alloc(III)* and *free(III)*. The function *xfreeall()* frees all memory allocated by *xalloc()* (it calls *brk(II)*). *Xalloc()* returns the address of the allocated area on success, and the value of *fatal()* on failure. *Xfree()* and *xfreeall()* don't return anything. *Xalloc()* uses a "first fit" strategy (unlike *alloc(II)*). *Xfree()* always coalesces contiguous free blocks. *Xalloc()* always allocates 2-byte words. *Xalloc()* actually allocates one more word than the amount requested. The extra word (the first word of the allocated block) contains the size (in bytes) of the entire block. This size is used by *xfree()* to identify contiguous blocks, and is used by *xalloc()* to implement the first fit strategy. Bad things will happen if that first (size) word is overwritten. Worse things happen if *xfree()* is called with a garbage argument.

xcreat(name, mode)

Xcreat() is used in the same way as *creat(II)*. *Xcreat()* requires write permission in the pertinent directory in *all* cases, and the created file is guaranteed to have the specified *mode* and be owned by the effective user (*xcreat()* does this by first unlinking the file to be created); *xcreat()* returns a file descriptor on success, and the value of *fatal()* on failure.

xfcreat(file, mode)

Xfcreat() is a macro that combines *xcreat()* and *fdlopen()*; its definition is:

```
fdlopen(xcreat(file, mode), 1)
```

xlopen(file, mode)

Xlopen() is a macro that combines *xopen()* (see below) and *fdlopen()*; its definition is:

```
fdlopen(xopen(file, mode), mode)
```

xlink(f1, f2)

Xlink() is used in the same way as *link(II)*. It is an interface to *link(II)* that handles all error conditions. It returns 0 on success, and the value of *fatal()* on failure.

xmsg(file, funcname)

Xmsg() is used by the other *x*-routines to generate an error message based on *errno* (see *INTRO(II)*). It calls *fatal()* with the appropriate error message. The second argument is a pointer to the calling function's name (a string). There are predefined messages for the most common errors. Other errors cause a message of the form:

str = sprintf(space, "error = %d, function = '%s'", errno, funcname)

to be passed to *fatal()*.

xopen(name, mode)

Xopen() is used in the same way as *open(II)*. It is an interface to *open(II)* that handles all error conditions. It returns a file descriptor on success, and the value of *fatal()* on failure.

xpipe(t)

Xpipe() is used in the same way as *pipe(II)*. It is an interface to *pipe(II)* that handles all error conditions. It returns 0 on success, and the value of *fatal()* on failure.

xunlink(f)

Xunlink() is used in the same way as *unlink(II)*. It is an interface to *unlink(II)* that handles all error conditions. It returns 0 on success, and the value of *fatal()* on failure.

2.4 Write Set

write(fildes, buffer, nbytes)


Write() is used in the same way as *write(II)*. It is an interface to *syswrite()* (see below) that handles all error conditions. It returns the number of bytes written on success, and the value of *fatal()* on failure.

syswrite(fildes, buffer, nbytes)

Syswrite() is identical to *write(II)*, except that the name *write* has been changed to *syswrite*.







Guide to IBM Remote Job Entry for PWB/UNIX Users

A. L. Sabsevitz

Bell Laboratories
Piscataway, New Jersey 08854

1. PREFACE

A set of background processes supports remote job entry (RJE) from a PWB/UNIX* computer to IBM System/360 and /370 host computers. "Hasp" is the common name used for the collection of programs and for the file organization that provides this facility; it allows PWB/UNIX to communicate with IBM's Job Entry Subsystem by mimicking an IBM 2770 remote station. The *PWB/UNIX User's Manual* page *hasp*(VIII) summarizes their design and operating procedures. That manual also contains a terse description of the *send*(1) command, which is the user's primary interface to RJE.¹ These are the definitive sources for information about RJE. Although the word "Hasp" may be used in this guide, it represents all IBM RJE subsystems of the PWB/UNIX system.

This guide is a tutorial overview of RJE.² It is addressed to the user who needs to know how to use the system, but does *not* need to know details of its implementation. The two following sections constitute an introduction to RJE.

2. PRELIMINARIES

To become a PWB/UNIX user, you must receive a login name that identifies you to the PWB/UNIX system. You should also get a copy of the *PWB/UNIX User's Manual*. This is a fairly complete description of the system and includes a section entitled "How to Get Started," which introduces you to PWB/UNIX; you should read that section before proceeding with this guide.

In order to begin using RJE, you need only become familiar with a subset of basic commands. You must understand the directory structure of the file system, and you should know something about the attributes of files: see *chdir*(1), *chmod*(1), *chown*(1), *cp*(1), *ln*(1), *ls*(1), *mkdir*(1), *mv*(1), *rm*(1). You must know how to enter, edit, and examine text files: see *cat*(1), *ed*(1), *pr*(1). You should know how to communicate with other users and with the system: see *mail*(1), *mesg*(1), *who*(1), *write*(1). And, finally, you might have to know how to describe your terminal to the system: see *ascii*(V), *stty*(1), *tabs*(1).

3. BASIC RJE

Let's suppose that you have used the editor, *ed*(1), to create a file "jobfile" that contains your control statements (JCL) and input data. This file should look exactly like a card deck, except that alphabetic characters, for convenience, may be in either upper or lower case. Here is an example:

```
% cat jobfile
//gener job (9999,r740),pgmname,class=x usr=(mylogin,myplace)
//step exec pgm=iebgener
//sysprint dd sysout=a
//sysin dd dummy
//sysut2 dd sysout=a
//sysut1 dd *
    first card of data
    :
    last card of data
/*
```

* UNIX is a Trademark of Bell Laboratories.

1. In this paper, RJE refers to the PWB/UNIX facilities provided by *hasp*(VIII), and *not* to the Remote Job Entry feature of IBM's HASP or JES2 subsystems.

2. The original versions of this manual and of RJE itself were written by T. G. Lyons.

To submit this job for execution, you must invoke the *send(I)* command:

```
% send jobfile
```

The system will reply:

```
10 cards.  
Queued as /usr/hasp/xmit311.
```

Note that *send* tells you how many cards it submitted and reports the position that your job has been assigned in the queue of all jobs waiting to be transmitted to the host system. Until the transmission of the job actually begins, you can prevent the job from being transmitted by doing a "chmod 0" on the queued file to make it unreadable. For our example, you could say: "chmod 0 /usr/hasp/xmit311".

When your job is accepted by the host system, a job number will be assigned to it, and an acknowledgement message will be generated. This indicates that your job has been scheduled on the host system. Later, after the job has executed, its output will be returned to the PWB/UNIX system. You will be notified automatically of both of these events: if you are logged in when RJE detects these events, and if you are permitting messages to be sent to your terminal (see *mesg(I)*), the following two messages will be sent to you (still using the example above) when the job is scheduled and when the output is returned, respectively:

```
Two bells  
$12.18.42 JOB 384 ON RM4.RD1 -- GENER  PGMRNAME  
Bell  
  
Two bells  
12:21:54 /a1/user/rje/prnt0 384.gener ready  
Bell
```

The job-acknowledgement message is passed on directly from the host system, as indicated by the fact that it appears in upper case. The output-ready message is generated by RJE and appears in lower case. Two bells, with an interval of one second between them, precede each message. They should be interpreted as a warning to stop typing on your terminal, so that the imminent message is not interspersed with your typing.

If you are not logged in when one of these events occurs, or if you do not allow messages to be sent to your terminal, then the notification will be posted to you via the *mail(I)* command. You can prevent messages directly by executing the *mesg(I)* command, or indirectly by executing another command, such as *pr(I)*, which prohibits messages for as long as it is active. You may inspect (by invoking the *mail* command) your *.mail* file at any time for messages that have been diverted. For this example, this might look as follows:

```
% mail  
  
From rje Mon Aug 1 12:20:36 1977  
$12.18.42 JOB 384 ON RM4.RD1 -- GENER  PGMRNAME  
  
From rje Mon Aug 1 12:21:55 1977  
12:21:54 /a1/user/rje/prnt0 384.gener ready  
  
Save?
```

Note that there may exist a discrepancy between the host and PWB/UNIX clocks.

The job-acknowledgement message performs two functions. First, it confirms the fact that your job has been scheduled for eventual execution. Second, it assigns a number to the job in such a way that the number and the name together will uniquely identify the job for some period of time.

The output-ready message provides the name of a PWB/UNIX file into which output has been written and identifies the job to which the output belongs (see *ls(I)*):

```
% ls -l prnt0  
-r--r-xr-- 1 rje      1184 Aug  1 12:21 prnt0
```

Note that rje retains ownership of the output and allows you only read access to it. It is intended that you will inspect the file, perhaps extract some information from it and then promptly delete it (see *rm(1)*):

```
% rm -f prnt0
```

The retention of machine-generated files, such as RJE output, is discouraged. It is your responsibility to remove files from your RJE directory. Files of RJE output may not exceed 256K bytes. In addition, **only** files of 64K bytes are guaranteed to be accepted in their entirety. Limits of 64K, 128K, or 192K may be automatically enforced if file space gets scarce. Output beyond the current limit will be discarded, with no provision for retrieving it. The user should also be aware of the fact that RJE attempts to keep roughly 1000 "blocks" free on any file system it uses. Warning messages or suspension of certain functions will occur as this limit is approached.

The most elementary way to examine your output is to *cat* it to your terminal. The Appendix shows the result of listing the output of our sample job in this way. Printouts are stored with standard tabs to conserve space, so you must ensure that the tabs are set on your terminal at every eighth column across the entire line; *tabs(1)* will do that for you.³ Because PWB/UNIX has no high-volume printing capability, you should route to the host's printer any large listings of which you desire a hard copy.

The structure of an output listing will generally conform to the following sequence:

```
HASP log
jcl information
data sets
HASP end
```

"Burst" pages are discarded. Single, double, and triple spacing is reflected in the output file, but other forms controls, such as the skip to the top of a new page, are suppressed. Page boundaries are indicated by the presence of a space *character* at the end of the last line of each page.

The big file scanner *bfs(1)* or the context editor *ed(1)* provide a more flexible method than *cat(1)* for examining printed output; *bfs* can handle files of any size and is more efficient than *ed* for scanning files.

RJE is also capable of receiving punched output as formatted files (see *ebcdic(V)*); this format allows an exact representation of an arbitrary card deck to be stored on the PWB/UNIX machine. However, there are few commands that can be used to manipulate EBCDIC files. You will probably want to route your punched output to one of the host's output devices.

4. SEND COMMAND

The *send(1)* command is capable of more general processing than has been indicated in the previous section. In the first place, it will concatenate a sequence of files to create a single job stream. This allows files of JCL and files of data to be maintained separately on the PWB/UNIX machine. In addition, it recognizes any line of an input file that begins with the character "~" as being a *control* line that can call for the inclusion, inside the current file, of some other file. This allows you to "send" a top-level skeleton that "pulls" in subordinate files as needed. Some of these may be "virtual" files that actually consist of the output of PWB/UNIX commands or Shell procedures. Furthermore, the *send* command is able to collect input directly from a terminal, and can be instructed to prompt for required information.

Each source of input can contain a format specification that determines such things as how to expand tabs and how long can an input line be. The manual page for *fspec(V)* explains how to define such formats. When properly instructed, *send* will also replace arbitrarily defined "keywords" by other text strings or by EBCDIC character codes. (These two substitution facilities are useful in other applications besides RJE; for that reason, *send* may be invoked under the name *gath* to produce standard output *without* submitting an RJE job.)

3. If your terminal doesn't have tabs, you should use the *stty(1)* command to cause PWB/UNIX to automatically convert tabs to spaces on output to the terminal.

Two aspects of *send* with which everyone should be acquainted are the ability to specify to which computer a job is to be submitted, and the ability to verify a job prior to submission. To run our sample job on a host machine known to RJE as "A", we would issue the command:

```
% send A jobfile
```

When no host is explicitly cited, *send* makes a reasonable choice.

To verify the text of a collected job stream, without actually submitting it, set the "-lq" flags:

```
% send -lq jobfile
```

The complete list of arguments and flags that control the execution of *send* can be found in *send(1)*.

5. JOB STREAM

It is assumed that the job stream submitted as the result of a single execution of *send* consists of a single *job*, i.e., the file that is queued for transmission should contain one JOB card near the beginning and no others. A priority control card may legitimately precede the JOB card. The JOB card must conform to the local installation's standard. At BISP, it has the following structure:

```
//name job (acct[,...]),pgmname[,keywds=?] [usr=...]
```

6. USER SPECIFICATION

The "usr=..." field is required if any print or punch output is to be delivered to the PWB/UNIX user.

```
usr=(login,place[,level][,retry])
```

where *login* is the PWB/UNIX login name of the user, *level* is the desired level of notification (see end of this section for an explanation), *retry* is a one-character code specifying the number of attempts to retransmit an entire job if the transmission to the host computer is interrupted by an unrecoverable error (default is three attempts; the digits "1" through "9" specify that number of retries; "0", "y", or "Y" invoke the default; any other entry in this field limits the number of attempts to one, i.e., no retry), and *place* is as follows:

- A. If *place* is the name of a *directory* (writable by others), then the output file is placed there as a unique *print* or *punch* file (up to 500 of each allowed). The mode of the file will be 454.
- B. If *place* is the name of an existing, writable (by others), non-executable (by others) file, then the output file replaces it. The mode of the file will be 454.
- C. If *place* is the name of a non-existent file in a writable (by others) directory, then the output file is placed there. The mode of the file will be 454.
- D. If *place* is the name of an executable (by others) file, then the RJE output is set up as standard input to *place*, and *place* is executed. Five string arguments are passed to *place*. For example, if *place* is a shell procedure, the following arguments are passed as \$1 ... \$5:
 1. Flag indicating whether file space is scarce in the file system where *place* resides. 0 indicates that space is *not* scarce, while 1 indicates that it is.
 2. Job name.
 3. Programmer's name.
 4. Job number.
 5. Login name from the "usr=..." specification.A ":" is passed if a value is not present.
- E. In all other cases, the output will be thrown away.

The *place* value must not be a full pathname, unless it refers to an executable file (see D above). For cases A, B, and C above (and case D, if a full pathname is not supplied), the name of the user's login directory will be used to form a full pathname.

The "usr=..." field may occur anywhere within the first 100 card images sent and within the first 200 output images received by the PWB/UNIX system. The only restrictions are:

- Column one must contain a "/" or a "*".
- "usr=..." must begin after column 4 and must be preceded by a space.

Therefore, the "usr=..." field may be placed on the JOB card, a comment card, passed as data, etc.

For redirection of output by the host, a "usr=..." card, if not already present, must be supplied by the user. This can be done by placing a job step that creates this card before your output steps.

Messages generated by RJE or passed on from the host are assigned a level of importance ranging from 1 to 9. The levels currently in use are:

- 3 transmittal assurance
- 5 job acknowledgement
- 6 output ready message
- 7 transmit format error

The optional "level" field of the "usr=..." specification must be a one- or two-digit code. A message from the host with importance "x" (where x comes from the above list) is compared with each of the two decimal digits "mw". If $x \geq w$ and if the user is logged in and is accepting messages, the message will be written to his or her terminal. Otherwise, if $x \geq m$, the message will be mailed to the user. In all other cases, the message will be discarded. The default "level field" is "54". You should specify level "1" if you want to receive complete notification, and level "59" to divert the last three messages in the above list to your mailbox.

7. CONTROL CARDS

A number of control cards are recognized by the host's HASP subsystem. Two are of particular interest to RJE users, because they control the disposition of output. These are the ROUTE and OUTPUT cards, and their use is illustrated below. If used, both should be inserted into a job stream *immediately* after the JOB card.

The ROUTE card can be used to direct the entire printed or punched output of a job to a specified destination. Two cards are required to direct both outputs:

```
/*route      punch      local
/*route      print      rmt55
```

The ROUTE card has a fixed format. "Print" or "punch" must begin in column 10, and the destination field in column 16.

The proper use of the OUTPUT card is a bit more complicated. It allows you to associate parameters with all SYSOUT data sets whose forms numbers match the one specified on the OUTPUT card. The forms number is fictitious and may consist of up to four characters. A copy count and destination are among the parameters that may be associated with SYSOUT data sets in this way:

```
//name job ...
/*output py d=rmt56
/*output abcd d=local,n=2
//step exec ...
//prt1 dd sysout=a
//prt2 dd sysout=(a,,abcd)
//prt3 dd sysout=(a,,py)
//pnch dd sysout=(b,,abcd)
```

In the above example, one copy of *prt1* would be directed to the default destination and one copy of *prt3* to *rmt56*. Two copies each would be made of *prt2* and *pnch*, and they would remain at the *local* site.

8. MONITORING RJE

RJE is designed to be an autonomous facility that does not require manual supervision. RJE is initiated by the PWB/UNIX operator after system "reboots" and continues in execution indefinitely. Experience has proved it to be reasonably robust, although it is vulnerable to system crashes and reconfigurations.

Users have a right to assume that, if the PWB/UNIX system is up for production use, RJE should also be up. This implies more than an ability to execute the *send(I)* command, which should be available at all times. It means that queued jobs should be submitted to the host for execution and their output returned to the PWB/UNIX system. If a user cannot obtain any throughput from RJE, the user should so advise the PWB/UNIX operators.

The *rjestat(I)* command, invoked without the "-" argument, will report the status of all RJE links for which a given PWB/UNIX system is configured. It may sometimes also print a message of the day from RJE.

```
% rjestat
```

```
15:12:24 RJE to B is operating normally.
```

```
15:12:25 A is not responding to RJE.
```

```
(8 files queued since 14:34:26)
```

A parenthetical statement, such as the last line above, will summarize any backlog of queued files waiting to be transmitted to the host machine. A backlog that persists for 20 minutes or more often is an indication that there exists a problem with the corresponding RJE link.

A host machine may be reported to be not responding to RJE because it is down, or because of its operator's failure to initialize the associated line, or because of a communications hardware failure.

Appendix—Sample Output Listing

% cat rje/prnt0
14.40.31 JOB 384 \$HASP373 GENER STARTED - INIT 26 - CLASS X - SYS RRMA
14.40.32 JOB 384 \$HASP395 GENER ENDED

----- JES2 JOB STATISTICS -----

1 AUG 77 JOB EXECUTION DATE

54 CARDS READ

76 SYSOUT PRINT RECORDS

0 SYSOUT PUNCH RECORDS

0.01 MINUTES EXECUTION TIME

1 //GENER JOB (9999,R740),PGMRNAME,CLASS=X JOB 384
*** USR=(MYLOGIN,MYPLACE)
2 //IEBGENER EXEC PGM=IEBGENER
3 //SYSPRINT DD DUMMY
4 //SYSIN DD DUMMY
5 //SYSUT2 DD SYSOUT=A
6 //SYSUT1 DD *
//

IEF236I ALLOC. FOR GENER IEBGENER
IEF237I DMY ALLOCATED TO SYSPRINT
IEF237I DMY ALLOCATED TO SYSIN
IEF237I JES ALLOCATED TO SYSUT2
IEF237I JES ALLOCATED TO SYSUT1
IEF142I GENER IEBGENER - STEP WAS EXECUTED - COND CODE 0000
IEF285I JES2.JOB0384.S00102 SYSOUT
IEF285I JES2.JOB0384.S10101 SYSIN
IEF373I STEP /IEBGENER/ START 77242.1440
IEF374I STEP /IEBGENER/ STOP 77242.1440 CPU 0MIN 00.13SEC SRB 0MIN 00.01SEC VIRT 36K SYS 188K

..... SERVICE UNITS=0000174 SERVICE RATE=0000268 SERVICE UNITS/SECOND
..... PERFORMANCE GROUP=005
..... EXCP COUNT BY UNIT ADDRESS
IEF375I JOB /GENER / START 77242.1440
IEF376I JOB /GENER / STOP 77242.1440 CPU 0MIN 00.13SEC SRB 0MIN 00.01SEC

..... SERVICE UNITS=0000174 SERVICE RATE=0000268 SERVICE UNITS/SECOND
..... APPROXIMATE PROCESSING TIME= .01 MINUTES
..... EXCPS=000000000
..... PROJECTED CHARGES= .01

first line of data

⋮

last line of data

•OS/VS2 REL 3.7 JES2• END JOBNAME=GENER BIN=R740 JOB #=384 PGMRNAME
•OS/VS2 REL 3.7 JES2• END JOBNAME=GENER BIN=R740 JOB #=384 PGMRNAME
•OS/VS2 REL 3.7 JES2• END JOBNAME=GENER BIN=R740 JOB #=384 PGMRNAME
% rm -f rje/prnt0



Guide to IBM Remote Job Entry for PWB/UNIX Users

ADDENDUM #1

UNIVAC RJE

I. PREFACE [1.] *

A set of PWB/UNIX † background processes supports remote job entry (RJE) from a PWB/UNIX computer to UNIVAC 1100-series host computers. "Uvac" is the common name used for the collection of programs and for the file organization that provides this facility.

II. BASIC RJE [3.]

Alphabetic characters, for convenience, may be in either upper or lower case. The master space "@*" may be represented as "****". Here is an example:

```
% cat jobfile
`run  echo,acct-no,project-id      usr= (mylogin,myplace)
`ed,i  .elt
      first card of data
      :
      last card of data
`fin
```

The system will reply:

```
8 cards.
Queued as /usr/uvac/xmit311.
```

The job acknowledgement messages are:

```
Two bells
$12.18.42 001.ECHO STARTED ACCT-NO
Bell

Two bells
12:21:54 /a1/user/rje/prnt0 .echo ready
Bell
```

To route your printout to the host's printer, insert the following control card into the run stream following the RUN card:

```
`sym  print$,pr
```

"Burst" pages are *not* discarded and the reception of punched output is *not* supported.

III. JOB STREAM [5.]

The RUN card must conform to the local installation's standard. In general, it has the following format:

```
@run  name,acct-no,project-id  [ . usr= ...]
```

IV. USER SPECIFICATION [6.]

The "usr= ..." field may occur anywhere within the first 100 card images sent and within the first 200 output images received by the PWB/UNIX system. The only restrictions are:

- Column 1 must contain a "@*" or a "****".
- "usr= ..." must begin after column 4 and must be preceded by a space.

Therefore, the "usr= ..." field may be placed on the RUN card, a message card, passed as data, etc.

* Numbers enclosed in curly braces are section numbers of *Guide to IBM Remote Job Entry for PWB/UNIX Users* by A. I. Sabsevitz, September 1977.

† Unix is a Trademark of Bell Laboratories.

Guide to IBM Remote Job Entry for PWB/UNIX Users

V. CONTROL CARDS [7.]

This section of the guide is not applicable to UNIVAC RJE.

VI. MONITORING RJE [8.]

The interactive status terminal capability of the *rjestat(I)* command is *not* implemented.



SCCS/PWB

User's Manual

L. E. Bonanni
A. L. Glasser

November 1977

SCCS/PWB
User's Manual

CONTENTS

1. INTRODUCTION	1
2. SCCS FOR BEGINNERS	2
2.1 Terminology 2	
2.2 Creating an SCCS File—The “admin” Command 2	
2.3 Retrieving a File—The “get” Command 2	
2.4 Recording Changes—The “delta” Command 3	
2.5 More about the “get” Command 4	
2.6 The “help” command 5	
3. HOW DELTAS ARE NUMBERED	5
4. SCCS COMMAND CONVENTIONS	6
5. SCCS COMMANDS	7
5.1 get 8	
5.2 delta 14	
5.3 admin 16	
5.4 prt 17	
5.5 help 18	
5.6 rmdel 18	
5.7 chghist 18	
5.8 what 19	
5.9 sccsdiff 19	
5.10 comb 19	
6. SCCS FILES	20
6.1 Protection 20	
6.2 Format 21	
6.3 Auditing 21	
REFERENCES	22

SCCS/PWB User's Manual

L. E. Bonanni

Bell Laboratories
Piscataway, New Jersey 08854

A. L. Glasser

Bell Laboratories
Holmdel, New Jersey 07733

ABSTRACT

The Source Code Control System (SCCS) is a system for controlling changes to files of text (typically, the source code and documentation of software systems). It provides facilities for storing, updating, and retrieving any version of a file of text, for controlling updating privileges to that file, for identifying the version of a retrieved file, and for recording who made each change, when and where it was made, and why. SCCS is a collection of programs that run under the PWB/UNIX* time-sharing system.

This document, together with the *PWB/UNIX User's Manual* [4], is a complete user's guide to Version 4 of SCCS, and supersedes all previous versions of the SCCS/PWB manual; it covers the following topics:

- How to get started with SCCS.
- The version numbering scheme.
- Basic information needed for day-to-day use of SCCS commands, including a discussion of the more useful arguments.
- Protection and auditing of SCCS files, including the differences between the use of SCCS by *individual* users on one hand, and *groups* of users on the other.

Neither the implementation of SCCS nor the installation procedure for SCCS are described here.

I. INTRODUCTION

The Source Code Control System (SCCS) is a collection of PWB/UNIX [1] commands that help individuals or projects control and account for changes to files of text (typically, the source code and documentation of software systems). It is convenient to conceive of SCCS as a custodian of files; it allows retrieval of particular versions of the files, administers changes to them, controls updating privileges to them, and records who made each change, when and where it was made, and why. This is important in environments in which programs and documentation undergo frequent changes (because of maintenance and/or enhancement work), inasmuch as it is sometimes desirable to regenerate the version of a program or document as it was before changes were applied to it. Obviously, this could be done by keeping copies (on paper or other media), but this quickly becomes unmanageable and wasteful as the number of programs and documents increases. SCCS provides an attractive solution because it stores on disk the original file and, whenever changes are made to it, stores only the *changes*; each set of changes is called a "delta."

This document, together with the *PWB/UNIX User's Manual* [4], is a complete user's guide to Version 4 of SCCS. This manual contains the following sections:

- *SCCS for Beginners*: How to make an SCCS file, how to update it, and how to retrieve a version thereof.
- *How Deltas Are Numbered*: How versions of SCCS files are numbered and named.
- *SCCS Command Conventions*: Conventions and rules generally applicable to all SCCS commands.
- *SCCS Commands*: Explanation of all SCCS commands, with discussions of the more useful arguments.

* UNIX is a Trademark of Bell Laboratories.

- *ScCs Files*: Protection, format, and auditing of SCCS files, including a discussion of the differences between using SCCS as an individual and using it as a member of a group or project. The role of a "project SCCS administrator" is introduced.

2. SCCS FOR BEGINNERS

It is assumed that the reader knows how to log onto a PWB/UNIX system, create files, and use the text editor [2,3]. A number of terminal-session fragments are presented below. All of them should be tried: the best way to learn SCCS is to use it.

To supplement the material in this manual, the detailed SCCS command descriptions (appearing in alphabetical order in Section 1 of [4]) should be consulted. Section 5 below contains a list of all the SCCS commands. For the time being, however, only basic concepts will be discussed.

2.1 Terminology

Each SCCS file is composed of one or more sets of changes applied to the null (empty) version of the file, with each set of changes usually depending on all previous sets. Each set of changes is called a "delta" and is assigned a name, called the *ScCs IDentification string* (SID), composed of at most four components, only the first two of which will concern us for now; these are the "release" and "level" numbers, separated by a period. Hence, the first delta is called "1.1", the second "1.2", the third "1.3", etc. The release number can also be changed (usually, this indicates a major change to the file) as discussed below.

Each delta of an SCCS file defines a particular version of the file. For example, delta 1.5 defines version 1.5 of the SCCS file, obtained by applying to the null (empty) version of the file the changes that constitute deltas 1.1, 1.2, etc., up to and including delta 1.5 itself, in that order.

2.2 Creating an SCCS File—The "admin" Command

Consider, for example, a file called "lang" that contains a list of programming languages:

```
c
pl/i
fortran
cobol
algol
```

We wish to give custody of this file to SCCS. The following *admin* command (which is used to *adminis-ter* SCCS files) creates an SCCS file and initializes delta 1.1 from the file "lang":

```
admin -ilang s.lang
```

All SCCS files *must* have names that begin with "s.", hence, "s.lang". The *-i* keyletter, together with its value "lang", indicates that *admin* is to create a new SCCS file and *initialize* it with the contents of the file "lang". This initial version is a set of changes applied to the null SCCS file; it is delta 1.1.

The *admin* command replies:

```
No id keywords (cm7)
```

This is a warning message (which may also be issued by other SCCS commands) that is to be ignored for the purposes of this section. Its significance is described in Section 5.1 below.

The file "lang" should be removed (because it can be easily reconstructed by using the *get* command, below):

```
rm lang
```

2.3 Retrieving a File—The "get" Command

The command:

```
get s.lang
```

causes the creation (retrieval) of the latest version of file "s.lang", and prints the following messages:

```
1.1
5 lines
No id keywords (cm7)
```

This means that *get* retrieved version 1.1 of the file, which is made up of 5 lines of text. The retrieved text is placed in a file whose name is formed by deleting the "s." prefix from the name of the SCCS file; hence, the file "lang" is created.

The above *get* command simply creates the file "lang" read-only, and keeps no information whatsoever regarding its creation. On the other hand, in order to be able to subsequently apply changes to an SCCS file with the *delta* command (see below), the *get* command must be informed of your intention to do so. This is done as follows:

```
get -e s.lang
```

The *-e* keyletter causes *get* to create a file "lang" for both reading and writing (so that it may be edited) and places certain information about the SCCS file in another new file, called the *p-file*, that will be read by the *delta* command. The *get* command prints the same messages as before, except that the warning message is not issued.

The file "lang" may now be changed, for example, by:

```
cd lang
27
$a
snobol
ratfor
.
w
41
q
```

2.4 Recording Changes—The "delta" Command

In order to record within the SCCS file the changes that have been applied to "lang", execute:

```
delta s.lang
```

Delta prompts with:

```
comments?
```

the response to which should be a description of why the changes were made; for example:

```
comments? added more languages
```

Delta then reads the *p-file*, and determines what changes were made to the file "lang". It does this by doing its own *get* to retrieve the original version, and by applying *diff(1)*¹ to the original version and the edited version.

When this process is complete, at which point the changes to "lang" have been stored in "s.lang", *delta* outputs:

```
No id keywords (cm7)
1.2
2 inserted
0 deleted
5 unchanged
```

The number "1.2" is the name of the delta just created, and the next three lines of output refer to the number of lines in the file "s.lang".

1. All references of the form *name(N)* refer to item *name* in section *N* of the *PWB/UNIX User's Manual* [4].

2.5 More about the "get" Command

As we have seen:

```
get s.lang
```

retrieves the latest version (now 1.2) of the file "s.lang". This is done by starting with the original version of the file and successively applying deltas (the changes) in order, until all have been applied.

For our example, the following commands are all equivalent:

```
get s.lang
```

```
get -r1 s.lang
```

```
get -r1.2 s.lang
```

The numbers following the `-r` keyletter are SIDs (see Section 2.1 above). Note that omitting the level number of the SID (as in the second example above) is equivalent to specifying the *highest* level number that exists within the specified release. Thus, the second command requests the retrieval of the latest version in release 1, namely 1.2. The third command specifically requests the retrieval of a particular version, in this case, also 1.2.

Whenever a truly major change is made to a file, the significance of that change is usually indicated by changing the *release* number (first component of the SID) of the delta being made. Since normal, automatic, numbering of deltas proceeds by incrementing the level number (second component of the SID), we must indicate to SCCS that we wish to change the release number. This is done with the `get` command:

```
get -e -r2 s.lang
```

Because release 2 does not exist, `get` retrieves the latest version *before* release 2; it also interprets this as a request to change the release number of the delta we wish to create to 2, thereby causing it to be named 2.1, rather than 1.3. This information is conveyed to `delta` via the *p-file*. `Get` then outputs:

```
1.2  
7 lines
```

indicating the retrieval of version 1.2. If the file is now edited, for example, by:

```
ed lang  
41  
/cobol/d  
w  
35  
q
```

and `delta` executed:

```
delta s.lang  
comments? deleted cobol from list of languages
```

we will see, by `delta`'s output, that version 2.1 is indeed created:

```
No id keywords (cm7)  
2.1  
0 inserted  
1 deleted  
6 unchanged
```

Deltas may now be created in release 2 (deltas 2.2, 2.3, etc.), or another new release may be created in a similar manner. This process may be continued as desired.

2.6 The "help" command

If the command:

```
get abc
```

is executed, the following message will be output:

```
ERROR [abc]: not an SCCS file (col)
```

The string "col" is a code for the diagnostic message, and may be used to obtain a fuller explanation of that message by use of the *help* command:

```
help col
```

This produces the following output:

```
col:
"not an SCCS file"
A file that you think is an SCCS file
does not begin with the characters "s."
```

Thus, *help* is a useful command to use whenever there is any doubt about the meaning of an SCCS message. Fuller explanations of almost all SCCS messages may be found in this manner.

3. HOW DELTAS ARE NUMBERED

It is convenient to conceive of the deltas applied to an SCCS file as the nodes of a tree, in which the root is the initial version of the file. The root delta (node) is normally named "1.1" and successor deltas (nodes) are named "1.2", "1.3", etc. The components of the names of the deltas are called the "release" and the "level" numbers, respectively. Thus, normal naming of successor deltas proceeds by incrementing the level number, which is performed automatically by SCCS whenever a delta is made. In addition, the user may wish to change the *release* number when making a delta, to indicate that a major change is being made. When this is done, the release number also applies to all successor deltas, unless specifically changed again. Thus, the evolution of a particular file may be represented as in Figure 1.




Figure 1. Evolution of an SCCS File

Such a structure may be termed the "trunk" of the SCCS tree. It represents the normal *sequential* development of an SCCS file, in which changes that are part of any given delta are dependent upon *all* the preceding deltas.

However, there are situations in which it is necessary to cause a *branching* in the tree, in that changes applied as part of a given delta are *not* dependent upon all previous deltas. As an example, consider a program which is in production use at version 1.3, and for which development work on release 2 is already in progress. Thus, release 2 may already have some deltas, precisely as shown in Figure 1. Assume that a production user reports a problem in version 1.3, and that the nature of the problem is such that it cannot wait to be repaired in release 2. The changes necessary to repair the trouble will be applied as a delta to version 1.3 (the version in production use). This creates a new version that will then be released to the user, but will *not* affect the changes being applied for release 2 (i.e., deltas 1.4, 2.1, 2.2, etc.).

The new delta is a node on a "branch" of the tree, and its name consists of *four* components, namely, the release and level numbers, as with trunk deltas, plus the "branch" and "sequence" numbers, as follows:

```
release.level.branch.sequence
```

The *branch* number is assigned to each branch that is a descendant of a particular trunk delta, with the first such branch being 1, the next one 2, and so on. The *sequence* number is assigned, in order, to each delta on a *particular branch*. Thus, 1.3.1.2 identifies the second delta of the first branch that derives from delta 1.3. This is shown in Figure 2.




Figure 2. Tree Structure with Branch Deltas

The concept of branching may be extended to any delta in the tree; the naming of the resulting deltas proceeds in the manner just illustrated.

Two observations are of importance with regard to naming deltas. First, the names of trunk deltas contain exactly two components, and the names of branch deltas contain exactly four components. Second, the first two components of the name of branch deltas are always those of the ancestral trunk delta, and the branch component is assigned in the order of creation of the branch, independently of its location relative to the trunk delta. Thus, a branch delta may always be identified as such from its name. Although the ancestral trunk delta may be identified from the branch delta's name, it is *not* possible to determine the *entire* path leading from the trunk delta to the branch delta. For example, if delta 1.3 has one branch emanating from it, all deltas on that branch will be named 1.3.1.*n*. If a delta on this branch then has another branch emanating from it, all deltas on the new branch will be named 1.3.2.*n* (see Figure 3). The only information that may be derived from the name of delta 1.3.2.2 is that it is the *chronologically* second delta on the *chronologically* second branch whose *trunk* ancestor is delta 1.3. In particular, it is *not* possible to determine from the name of delta 1.3.2.2 all of the deltas between it and its trunk ancestor (1.3).




Figure 3. Extending the Branching Concept

It is obvious that the concept of branch deltas allows the generation of arbitrarily complex tree structures. Although this capability has been provided for certain specialized uses, it is strongly recommended that the SCCS tree be kept as simple as possible, because comprehension of its structure becomes extremely difficult as the tree becomes more complex.

4. SCCS COMMAND CONVENTIONS

This section discusses the conventions and rules that apply to SCCS commands. These rules and conventions are generally applicable to *all* SCCS commands, except as indicated below. SCCS commands accept two types of arguments: *keyletter* arguments and *file* arguments.

Keyletter arguments (hereafter called simply “keyletters”) begin with a minus sign (—), followed by a lower-case alphabetic character, and, in some cases, followed by a value. These keyletters control the execution of the command to which they are supplied.

File arguments (which may be names of files and/or directories) specify the file(s) that the given SCCS command is to process; naming a directory is equivalent to naming *all* the SCCS files within the directory. Non-SCCS files and unreadable² files in the named directories are silently ignored.

In general, file arguments may *not* begin with a minus sign. However, if the name “—” (a lone minus sign) is specified as an argument to a command, the command reads the standard input for lines and takes each line as the *name* of an SCCS file to be processed. The standard input is read until end-of-file. This feature is often used in pipelines [4] with, for example, the *find(1)* or *ls(1)* commands. Again, names of non-SCCS files and of unreadable files are silently ignored.

All keyletters specified for a given command apply to *all* file arguments of that command. All keyletters are processed before any file arguments, with the result that the placement of keyletters is arbitrary (i.e., keyletters may be interspersed with file arguments). File arguments, however, are processed left to right.

Somewhat different argument conventions apply to the *help*, *what*, and *sccsdiff* commands (see Sections 5.5, 5.8, and 5.9).

Certain actions of various SCCS commands are controlled by *flags* appearing in SCCS files. Some of these flags are discussed below. For a complete description of all such flags, see *admin(1)*.

The distinction between the *real user* and the *effective user* of a PWB/UNIX system is of concern in discussing various actions of SCCS commands. For the present, it is assumed that both the real user and the effective user are one and the same (i.e., the user who is logged into a PWB/UNIX system); this subject is further discussed in Section 6.1.

All SCCS commands that modify an SCCS file do so by writing a temporary copy, called the *x-file*, which ensures that the SCCS file will not be damaged should processing terminate abnormally. The name of the *x-file* is formed by replacing the “s.” of the SCCS file name with “x.”. When processing is complete, the old SCCS file is removed and the *x-file* is renamed to be the SCCS file. The *x-file* is created in the directory containing the SCCS file, is given the same mode (see *chmod(1)*) as the SCCS file, and is owned by the effective user.

To prevent simultaneous updates to an SCCS file, commands that modify SCCS files create a *lock-file*, called the *z-file*, whose name is formed by replacing the “s.” of the SCCS file name with “z.”. The *z-file* contains the *process number* [1] of the command that creates it, and its existence is an indication to other commands that that SCCS file is being updated. Thus, other commands that modify SCCS files will not process an SCCS file if the corresponding *z-file* exists. The *z-file* is created with mode 444 (read-only) in the directory containing the SCCS file, and is owned by the effective user. This file exists only for the duration of the execution of the command that creates it. In general, users can ignore *x-files* and *z-files*; they may be useful in the event of system crashes or similar situations.

SCCS commands produce diagnostics (on the diagnostic output [5]) of the form:

ERROR [name-of-file-being-processed]: message text (code)

The *code* in parentheses may be used as an argument to the *help* command (see Section 5.5) to obtain a further explanation of the diagnostic message.

Detection of a fatal error during the processing of a file causes the SCCS command to terminate processing of *that* file and to proceed with the next file, in order, if more than one file has been named.

5. SCCS COMMANDS

This section describes the major features of all the SCCS commands. Detailed descriptions of the commands and of all their arguments are given in [4], and should be consulted for further information. The discussion below covers only the more common arguments of the various SCCS commands.

2. Because of permission modes (see *chmod(1)*).

Because the commands *get* and *delta* are the most frequently used, they are presented first. The other commands follow in approximate order of importance.

The following is a summary of all the SCCS commands and of their major functions:

<i>get</i>	Retrieves versions of SCCS files.
<i>delta</i>	Applies changes (deltas) to the text of SCCS files, i.e., creates new versions.
<i>admin</i>	Creates SCCS files and applies changes to parameters of SCCS files.
<i>prt</i>	Formats and prints portions of SCCS files.
<i>help</i>	Gives explanations of diagnostic messages.
<i>rmdel</i>	Removes a delta from an SCCS file; allows the removal of deltas that were created by mistake.
<i>chghist</i>	Changes the commentary associated with a delta.
<i>what</i>	Searches any PWB/UNIX file(s) for all occurrences of a special pattern and prints out what follows it; is useful in finding identifying information inserted by the <i>get</i> command.
<i>sccsdiff</i>	Shows the differences between any two versions of an SCCS file.
<i>comb</i>	Combines two or more consecutive deltas of an SCCS file into a single delta; often reduces the size of the SCCS file.

5.1 *get*

The *get* command creates a text file that contains a particular version of an SCCS file. The particular version is retrieved by beginning with the initial version, and then applying deltas, in order, until the desired version is obtained. The created file is called the *g-file*; its name is formed by removing the "s." from the SCCS file name. The *g-file* is created in the current directory [1] and is owned by the real user. The mode assigned to the *g-file* depends on how the *get* command is invoked, as discussed below.

The most common invocation of *get* is:

```
get s.abc
```

which normally retrieves the latest version on the trunk of the SCCS file tree, and produces (for example) on the standard output [5]:

```
1.3  
67 lines  
No id keywords (cm7)
```

which indicates that:

1. Version 1.3 of file "s.abc" was retrieved (1.3 is the latest trunk delta).
2. This version has 67 lines of text.
3. No ID keywords were substituted in the file (see Section 5.1.1 for a discussion of ID keywords).

The generated *g-file* (file "abc") is given mode .444 (read-only), since this particular way of invoking *get* is intended to produce *g-files* only for inspection, compilation, etc., and *not* for editing (i.e., *not* for making deltas).

In the case of several file arguments (or directory-name arguments), similar information is given for each file processed, but the SCCS file name precedes it. For example:

```
get s.abc s.def
```

produces:

```
s.abc:
1.3
67 lines
No id keywords (cm7)
```

```
s.def:
1.7
85 lines
No id keywords (cm7)
```

5.1.1 *Id Keywords*

In generating a *g-file* to be used for compilation, it is useful and informative to record the date and time of creation, the version retrieved, the module's name, etc., within the *g-file*, so as to have this information appear in a load module when one is eventually created. SCCS provides a convenient mechanism for doing this automatically. *Identification (ID) Keywords* appearing anywhere in the generated file are replaced by appropriate values according to the definitions of these ID keywords. The format of an ID keyword is an upper-case letter enclosed by percent signs (%). For example:

```
%I%
```

is defined as the ID keyword that is replaced by the SID of the retrieved version of a file. Similarly, %H% is defined as the ID keyword for the current date (in the form "mm/dd/yy"), and %M% is defined as the name of the *g-file*. Thus, executing *get* on an SCCS file that contains the PL/I declaration:

```
DCL ID CHAR(100) VAR INIT('%M% %I% %H%');
```

gives (for example) the following:

```
DCL ID CHAR(100) VAR INIT('MODNAME 2.3 07/07/77');
```

When no ID keywords are substituted by *get*, the following message is issued:

```
No id keywords (cm7)
```

This message is normally treated as a warning by *get*, although the presence of the *i* flag in the SCCS file causes it to be treated as an error (see Section 5.2 for further information).

For a complete list of the approximately twenty ID keywords provided, see *get(I)*.

5.1.2 *Retrieval of Different Versions*

Various keyletters are provided to allow the retrieval of other than the default version of an SCCS file. Normally, the default version is the most recent delta of the highest-numbered release on the *trunk* of the SCCS file tree. However, if the SCCS file being processed has a *d* (default SID) flag, the SID specified as the value of this flag is used as a default. The default SID is interpreted in exactly the same way as the value supplied with the *-r* keyletter of *get*.

The *-r* keyletter is used to specify an SID to be retrieved, in which case the *d* (default SID) flag (if any) is ignored. For example:

```
get -r1.3 s.abc
```

retrieves version 1.3 of file "s.abc", and produces (for example) on the standard output:

```
1.3
64 lines
```

A branch delta may be retrieved similarly:

```
get -r1.5.2.3 s.abc
```

which produces (for example) on the standard output:

```
1.5.2.3
234 lines
```

When a two- or four-component SID is specified as a value for the *-r* keyletter (as above) and the

particular version does not exist in the sccs file, an error message results. Omission of the level number, as in:

```
get -r3 s.abc
```

causes retrieval of the *trunk* delta with the highest level number within the given release, if the given release exists. Thus, the above command might output:

```
3.7  
213 lines
```

If the given release does not exist, *get* retrieves the *trunk* delta with the highest level number within the highest-numbered existing release that is lower than the given release. For example, assuming release 9 does not exist in file "s.abc", and that release 7 is actually the highest-numbered release below 9, execution of:

```
get -r9 s.abc
```

might produce:

```
7.6  
420 lines
```

which indicates that trunk delta 7.6 is the latest version of file "s.abc" below release 9. Similarly, omission of the sequence number, as in:

```
get -r4.3.2 s.abc
```

results in the retrieval of the branch delta with the highest sequence number on the given branch, if it exists. (If the given branch does not exist, an error message results.) This might result in the following output:

```
4.3.2.8  
89 lines
```

The *-t* keyletter is used to retrieve the latest ("top") version in a particular *release* (i.e., when no *-r* keyletter is supplied, or when its value is simply a release number). The latest version is defined as that delta which was produced most recently, independent of its location on the sccs file tree. Thus, if the most recent delta in release 3 is 3.5,

```
get -r3 -t s.abc
```

might produce:

```
3.5  
59 lines
```

However, if branch delta 3.2.1.5 were the latest delta (created after delta 3.5), the same command might produce:

```
3.2.1.5  
46 lines
```

5.1.3 Retrieval with Intent to Make a Delta

Specification of the *-e* keyletter to the *get* command is an indication of the intent to make a delta, and, as such, its use is restricted. The presence of this keyletter causes *get* to:

1. Check the *user list* (which is the list of *login* names of users allowed to make deltas (see Section 6.2)) to determine if the *login* name of the user executing *get* is on that list. Note that a *null* (empty) user list behaves as if it contained *all* possible *login* names.
2. Check that the *release* (R) of the version being retrieved satisfies the relation:

$$\text{floor} \leq R \leq \text{ceiling}$$

to determine if the release being accessed is a protected release. The *floor* and *ceiling* are specified as *flags* in the sccs file.

A failure of either condition causes the processing of that SCCS file to terminate.

If the above checks succeed, the `-e` keyletter causes the creation of a *g-file* in the current directory with mode 644 (readable by everyone, writable only by the owner) owned by the real user. If a *writable g-file* already exists, *get* terminates with an error. This is to prevent inadvertent destruction of a *g-file* that already exists and is being edited for the purpose of making a delta.

Any ID keywords appearing in the *g-file* are *not* substituted by *get* when the `-e` keyletter is specified, because the generated *g-file* is to be subsequently used to create another delta, and replacement of ID keywords would cause them to be permanently changed within the SCCS file. In view of this, *get* does not need to check for the presence of ID keywords within the *g-file*, so that the message:

```
No id keywords (cm7)
```

is never output when *get* is invoked with the `-e` keyletter.

In addition, the `-e` keyletter causes the creation (or updating) of a *p-file*, which is used to pass information to the *delta* command (see Section 5.1.4).

The following is an example of the use of the `-e` keyletter:

```
get -e s.abc
```

which produces (for example) on the standard output:

```
1.3
67 lines
```

If the `-r` and/or `-t` keyletters are used together with the `-e` keyletter, the version retrieved for editing is as specified by the `-r` and/or `-t` keyletters.

The keyletters `-i` and `-x` may be used to specify a list (see *get(I)* for the syntax of such a list) of deltas to be *included* and *excluded*, respectively, by *get*. Including a delta means forcing the changes that constitute the particular delta to be included in the retrieved version. This is useful if one wants to apply the same changes to more than one version of the SCCS file. Excluding a delta means forcing it to be *not* applied. This may be used to undo, in the version of the SCCS file to be created, the effects of a previous delta. Whenever deltas are included or excluded, *get* checks for possible interference between such deltas and those deltas that are normally used in retrieving the particular version of the SCCS file. (Two deltas can interfere, for example, when each one changes the same line of the retrieved *g-file*.) Any interference is indicated by a warning that shows the range of lines within the retrieved *g-file* in which the problem may exist. The user is expected to examine the *g-file* to determine whether a problem actually exists, and to take whatever corrective measures (if any) are deemed necessary (e.g., edit the file).

■ *The `-i` and `-x` keyletters should be used with extreme care.*

The `-k` keyletter is provided to facilitate regeneration of a *g-file* that may have been accidentally removed or ruined subsequent to the execution of *get* with the `-e` keyletter, or to simply generate a *g-file* in which the replacement of ID keywords has been suppressed. Thus, a *g-file* generated by the `-k` keyletter is identical to one produced by *get* executed with the `-e` keyletter. However, no processing related to the *p-file* takes place.

5.1.4 The *p-file* and Concurrent Deltas

The ability to retrieve different versions of an SCCS file allows a number of deltas to be “in progress” at any given time. This means that a number of *get* commands with the `-e` keyletter may be executed on the same file, provided that no two executions retrieve the same version *nor* lead to the subsequent creation of the same version by *delta*.

The *p-file* (which is created by the *get* command invoked with the `-e` keyletter) is named by replacing the “s.” in the SCCS file name with “p.”. It is created in the directory containing the SCCS file, is given mode 644 (readable by everyone, writable only by the owner), and is owned by the effective user. The *p-file* contains the following information for each delta that is still “in progress”:³

³ Other information may be present, but is not of concern here. See *get(I)* for further discussion.

- The SID of the retrieved version.
- The SID that will be given to the new delta when it is created.
- The login name of the real user executing *get*.

The first execution of “*get -e*” causes the *creation* of the *p-file* for the corresponding SCCS file. Subsequent executions only *update* the *p-file* by inserting a line containing the above information. Before inserting this line, however, *get* checks that:

- No entry already in the *p-file* specifies as already retrieved the SID of the version to be retrieved.
- That the new (“to-be-created”) SID is not already specified as such in the *p-file*.

If both checks succeed, the user is informed that other deltas are in progress, and processing continues. If either check fails, an error message results. It is important to note that the various executions of *get* should be carried out from different directories. Otherwise, only the first execution will succeed, since subsequent executions would attempt to over-write a *writable g-file*, which is an SCCS error condition. In practice, such multiple executions are performed by different users,⁴ so that this problem does not arise, since each user normally has a different working directory [5].

Table 1 shows, for the most useful cases, what version of an SCCS file is retrieved by *get*, as well as the SID of the version to be eventually created by *delta*, as a function of the SID specified to *get*.

5.1.5 Keyletters That Affect Output

Specification of the **-p** keyletter causes *get* to write the retrieved text to the standard output, rather than to a *g-file*. In addition, all output normally directed to the standard output (such as the SID of the version retrieved and the number of lines retrieved) is directed instead to the diagnostic output. This may be used, for example, to create *g-files* with arbitrary names:

```
get -p s.abc > arbitrary-filename
```

The **-p** keyletter is particularly useful when used with the “!” or “\$” arguments of the PWB/UNIX *send(I)* command. For example:

```
send MOD=s.abc REL=3 compile
```

if file “compile” contains:

```
//plicomp job job-card-information
//step1 exec plicke
//pli.sysin dd *
~s
~!get -p -rREL MOD
/*
//
```

will *send* the highest level of release 3 of file “s.abc”. Note that the line “~s”, which causes *send(I)* to make ID keyword substitutions before detecting and interpreting control lines, is necessary if *send(I)* is to substitute “s.abc” for MOD and “3” for REL in the line “~!get -p -rREL MOD”.

The **-s** keyletter suppresses all output that is *normally* directed to the standard output. Thus, the SID of the retrieved version, the number of lines retrieved, etc., are not output. This does not, however, affect messages to the diagnostic output. This keyletter is used to prevent non-diagnostic messages from appearing on the user’s terminal, and is often used in conjunction with the **-p** keyletter to “pipe” the output of *get*, as in:

```
get -p -s s.abc | nroff
```

The **-g** keyletter is supplied to suppress the actual retrieval of the text of a version of the SCCS file. This may be useful in a number of ways. For example, to verify the existence of a particular SID in an SCCS file, one may execute:

```
get -g -r4.3 s.abc
```

4. See Section 6.1 for a discussion of how different users are permitted to use SCCS commands on the same files.

TABLE 1. Determination of New SID

Case	SID Specified*	-b Keyletter Used†	Other Conditions	SID Retrieved	SID of Delta to be Created
1.	none‡	no	R defaults to mR	mR.mL	mR.(mL + 1)
2.	none‡	yes	R defaults to mR	mR.mL	mR.mL.(mB + 1).1
3.	R	no	R > mR	mR.mL	R.1§
4.	R	no	R = mR	mR.mL	mR.(mL + 1)
5.	R	yes	R > mR	mR.mL	mR.mL.(mB + 1).1
6.	R	yes	R = mR	mR.mL	mR.mL.(mB + 1).1
7.	R	—	R < mR and R does <i>not</i> exist	hR.mL**	hR.mL.(mB + 1).1
8.	R	—	Trunk successor in release > R and R exists	R.mL	R.mL.(mB + 1).1
9.	R.L	no	No trunk successor	R.L	R.(L + 1)
10.	R.L	yes	No trunk successor	R.L	R.L.(mB + 1).1
11.	R.L	—	Trunk successor in release ≥ R	R.L	R.L.(mB + 1).1
12.	R.L.B	no	No branch successor	R.L.B.mS	R.L.B.(mS + 1)
13.	R.L.B	yes	No branch successor	R.L.B.mS	R.L.(mB + 1).1
14.	R.L.B.S	no	No branch successor	R.L.B.S	R.L.B.(S + 1)
15.	R.L.B.S	yes	No branch successor	R.L.B.S	R.L.(mB + 1).1
16.	R.L.B.S	—	Branch successor	R.L.B.S	R.L.(mB + 1).1

* "R", "L", "B", and "S" are the "release", "level", "branch", and "sequence" components of the SID, respectively; "m" means "maximum". Thus, for example, "R.mL" means "the maximum level number within release R"; "R.L.(mB + 1).1" means "the first sequence number on the *new* branch (i.e., maximum branch number plus 1) of level L within release R". Note that if the SID specified is of the form "R.L", "R.L.B", or "R.L.B.S", each of the specified components *must* exist.

† The -b keyletter is effective only if the b flag (see *admin(1)*) is present in the file. In this table, an entry of "—" means "irrelevant".

‡ This case applies if the d (default SID) flag is *not* present in the file. If the d flag *is* present in the file, then the SID obtained from the d flag is interpreted as if it had been specified on the command line. Thus, one of the other cases in this table applies.

§ This case is used to force the creation of the *first* delta in a *new* release.

** "hR" is the highest *existing* release that is lower than the specified, *nonexistent*, release R.

This outputs the given SID if it exists in the SCCS file, or it generates an error message, if it does not. Another use of the -g keyletter is in regenerating a *p-file* that may have been accidentally destroyed:

```
get -c -g s.abc
```

The -l keyletter causes the creation of an *l-file*, which is named by replacing the "s." of the SCCS file name with "l.". This file is created in the current directory, with mode 444 (read-only), and is owned by the real user. It contains a table (whose format is described in *get(1)*) showing which deltas were used in constructing a particular version of the SCCS file. For example:

```
get -r2.3 -l s.abc
```

generates an *l-file* showing which deltas were applied to retrieve version 2.3 of the SCCS file. Specifying a *value* of "p" with the -l keyletter, as in:

```
get -lp -r2.3 s.abc
```

causes the generated output to be written to the standard output rather than to the *l-file*. Note that the -g keyletter may be used with the -l keyletter to suppress the actual retrieval of the text.

The `-m` keyletter is of use in identifying, line by line, the changes applied to an SCCS file. Specification of this keyletter causes each line of the generated *g-file* to be preceded by the SID of the delta that caused that line to be inserted. The SID is separated from the text of the line by a tab character.

The `-n` keyletter causes each line of the generated *g-file* to be preceded by the value of the `%M%` ID keyword (see Section 5.1.1) and a tab character. The `-n` keyletter is most often used in a pipeline with `grep(1)`. For example, to find, in the latest version of each SCCS file in a directory, all lines that match a given pattern, the following may be executed:

```
get -p -n -s directory | grep pattern
```

If both the `-m` and `-n` keyletters are specified, each line of the generated *g-file* is preceded by the value of the `%M%` ID keyword and a tab (this is the effect of the `-n` keyletter), followed by the line in the format produced by the `-m` keyletter. Because use of the `-m` keyletter and/or the `-n` keyletter causes the contents of the *g-file* to be modified, such a *g-file* must *not* be used for creating a delta. Therefore, neither the `-m` keyletter nor the `-n` keyletter may be specified together with the `-e` keyletter.

See `get(1)` for a full description of additional `get` keyletters.

5.2 delta

The `delta` command is used to incorporate the changes made to a *g-file* into the corresponding SCCS file, i.e., to create a delta, and, therefore, a new version of the file.

Invocation of the `delta` command requires the existence of a *p-file* (see Sections 5.1.3 and 5.1.4). `Delta` examines the *p-file* to verify the presence of an entry containing the user's login name. If none is found, an error message results. `Delta` also performs the same permission checks that `get` performs when invoked with the `-e` keyletter. If all checks are successful, `delta` determines what has been changed in the *g-file*, by comparing it (via `diff(1)`) with its own, temporary copy of the *g-file* as it was before editing. This temporary copy of the *g-file* is called the *d-file* (its name is formed by replacing the "s." of the SCCS file name with "d.") and is obtained by performing an internal `get` at the SID specified in the *p-file* entry.

The required *p-file* entry is the one containing the login name of the user executing `delta`, because the user who retrieved the *g-file* must be the one who will create the delta. However, if the login name of the user appears in more than one entry (i.e., the same user executed `get` with the `-e` keyletter more than once on the same SCCS file), the `-r` keyletter must be used with `delta` to specify the SID that is to be used by the internal `get` to obtain the *d-file*. The SID specified must, of course, appear in one of the entries in the *p-file*; this entry is the one used to obtain the SID of the delta to be created.

In practice, the most common invocation of `delta` is:

```
delta s.abc
```

which prompts on the standard output (but only if it is a terminal):

```
comments?
```

to which the user replies with a description of why the delta is being made, terminating the reply with a newline character. The user's response may be up to 512 characters long, with newlines *not* intended to terminate the response escaped by "\".

If the SCCS file has a `v` flag, `delta` first prompts with:

```
MRs?
```

on the standard output. (Again, this prompt is printed only if the standard output is a terminal.) The standard input is then read for MR⁵ numbers, separated by blanks and/or tabs, terminated in the same manner as the response to the prompt "comments?".

5. In a tightly controlled environment, it is expected that deltas are created only as a result of some trouble report, change request, trouble ticket, etc. (collectively called here Modification Requests, or MRs) and that it is desirable or necessary to record such MR number(s) within each delta.

The `-y` and/or `-m` keyletters are used to supply the commentary (comments and MR numbers, respectively) on the command line, rather than through the standard input. For example:

```
delta -y"descriptive comment" -m"mrnum1 mrnum2" s.abc
```

In this case, the corresponding prompts are not printed, and the standard input is not read. The `-m` keyletter is allowed only if the SCCS file has a `v` flag. These keyletters are useful when *delta* is executed from within a *Shell procedure* (see *sh(1)*).

The commentary (comments and/or MR numbers), whether solicited by *delta* or supplied via keyletters, is recorded as part of the entry for the delta being created, and applies to *all* SCCS files processed by the same invocation of *delta*. This implies that if *delta* is invoked with more than one file argument, and the first file named has a `v` flag, all files named must have this flag. Similarly, if the first file named does not have this flag, then none of the files named may have it. Any file that does not conform to these rules is not processed.

When processing is complete, *delta* outputs (on the standard output) the SID of the created delta (obtained from the *p-file* entry) and the counts of lines inserted, deleted, and left unchanged by the delta. Thus, a typical output might be:

```
1.4
14 inserted
7 deleted
345 unchanged
```

It is possible that the counts of lines reported as inserted, deleted, or unchanged by *delta* do not agree with the user's perception of the changes applied to the *g-file*. The reason for this is that there usually are a number of ways to describe a set of such changes, especially if lines are moved around in the *g-file*, and *delta* is likely to find a description that differs from the user's perception. However, the *total* number of lines of the new delta (the number inserted plus the number left unchanged) should agree with the number of lines in the edited *g-file*.

If, in the process of making a delta, *delta* finds no ID keywords in the edited *g-file*, the message:

```
No id keywords (cm7)
```

is issued after the prompts for commentary, but before any other output. This indicates that any ID keywords that may have existed in the SCCS file have been replaced by their values, or deleted during the editing process. This could be caused by creating a delta from a *g-file* that was created by a *get* without the `-e` keyletter (recall that ID keywords are replaced by *get* in that case), or by accidentally deleting or changing the ID keywords during the editing of the *g-file*. Another possibility is that the file may never have had any ID keywords. In any case, it is left up to the user to determine what remedial action is necessary, but the delta is made, unless there is an `i` flag in the SCCS file, indicating that this should be treated as a fatal error. In this last case, the delta is not created.

After processing of an SCCS file is complete, the corresponding *p-file* entry is removed from the *p-file*.⁶ If there is only *one* entry in the *p-file*, then the *p-file* itself is removed.

In addition, *delta* removes the edited *g-file*, unless the `-n` keyletter is specified. Thus:

```
delta -n s.abc
```

will keep the *g-file* upon completion of processing.

The `-s` ("silent") keyletter suppresses all output that is normally directed to the standard output, other than the prompts "comments?" and "MRs?". Thus, use of the `-s` keyletter together with the `-y` keyletter (and possibly, the `-m` keyletter) causes *delta* neither to read the standard input nor to write the standard output.

6. All updates to the *p-file* are made to a temporary copy, the *q-file*, whose use is similar to the use of the *x-file*, which is described in Section 4 above.

The differences between the *g-file* and the *d-file* (see above), which constitute the delta, may be printed on the standard output by using the `-p` keyletter. The format of this output is similar to that produced by `diff(1)`.

5.3 admin

The `admin` command is used to *administer* SCCS files, that is, to create new SCCS files and to change parameters of existing ones. When an SCCS file is created, its parameters are initialized by use of keyletters or are assigned default values if no keyletters are supplied. The same keyletters are used to change the parameters of existing files.

Two keyletters are supplied for use in conjunction with detecting and correcting "corrupted" SCCS files, and are discussed in Section 6.3 below.

Newly-created SCCS files are given mode 444 (read-only) and are owned by the effective user.

Only a user with write permission in the directory containing the SCCS file may use the `admin` command upon that file.

5.3.1 Creation of SCCS Files

An SCCS file may be created by executing the command:

```
admin -ifirst s.abc
```

in which the value ("first") of the `-i` keyletter specifies the name of a file from which the text of the *initial* delta of the SCCS file "s.abc" is to be taken. Omission of the value of the `-i` keyletter indicates that `admin` is to read the standard input for the text of the initial delta. Thus, the command:

```
admin -i s.abc < first
```

is equivalent to the previous example. If the text of the initial delta does not contain ID keywords, the message:

```
No id keywords (cm7)
```

is issued by `admin` as a warning. However, if the same invocation of the command also sets the `i` flag (not to be confused with the `-i` keyletter), the message is treated as an error and the SCCS file is not created. Only *one* SCCS file may be created at a time using the `-i` keyletter.

When an SCCS file is created, the *release* number assigned to its first delta is normally "1", and its *level* number is always "1". Thus, the first delta of an SCCS file is normally "1.1". The `-r` keyletter is used to specify the release number to be assigned to the first delta. Thus:

```
admin -ifirst -r3 s.abc
```

indicates that the first delta should be named "3.1" rather than "1.1". Because this keyletter is only meaningful in creating the first delta, its use is only permitted with the `-i` keyletter.

5.3.2 Initialization and Modification of SCCS File Parameters

The portion of the SCCS file reserved for *descriptive text* (see Section 6.2) may be initialized or changed through the use of the `-t` keyletter. The descriptive text is intended as a summary of the contents and purpose of the SCCS file, although its contents may be arbitrary, and it may be arbitrarily long.

When an SCCS file is being created and the `-t` keyletter is supplied, it must be followed by the name of a file from which the descriptive text is to be taken. For example, the command:

```
admin -ifirst -tdesc s.abc
```

specifies that the descriptive text is to be taken from file "desc".

When processing an *existing* SCCS file, the `-t` keyletter specifies that the descriptive text (if any) currently in the file is to be *replaced* with the text in the named file. Thus:

```
admin -tdesc s.abc
```

specifies that the descriptive text of the SCCS file is to be replaced by the contents of “desc”; omission of the file name after the **-t** keyletter as in:

```
admin -t s.abc
```

causes the *removal* of the descriptive text from the SCCS file.

The *flags* (see Section 6.2) of an SCCS file may be initialized, changed, or deleted through the use of the **-f** and **-d** keyletters, respectively. The flags of an SCCS file are used to direct certain actions of the various commands. See *admin(1)* for a description of all the flags. For example, the **v** flag specifies that *delta* is to prompt for Modification Request (MR) numbers, and the **d** (default *sid*) flag specifies the default version of the SCCS file to be retrieved by the *get* command. The **-f** keyletter is used to set a flag and, possibly, to set its value. For example:

```
admin -ifirst -fv -fmmodname s.abc
```

sets the **v** flag and the **m** (module name) flag. The value “modname” specified for the **m** flag is the value that the *get* command will use to replace the **%M%** keyword. (In the absence of the **m** flag, the name of the *g-file* is used as the replacement for the **%M%** keyword.) Note that several **-f** keyletters may be supplied on a single invocation of *admin*, and that **-f** keyletters may be supplied whether the command is creating a new SCCS file or processing an existing one.

The **-d** keyletter is used to delete a flag from an SCCS file, and may only be specified when processing an existing file. As an example, the command:

```
admin -dm s.abc
```

removes the **m** flag from the SCCS file. Several **-d** keyletters may be supplied on a single invocation of *admin*, and may be intermixed with **-f** keyletters.

SCCS files contain a list (*user list*) of login names of users who are allowed to create deltas (see Sections 5.1.3 and 6.2). This list is empty by default, which implies that *anyone* may create deltas. To add login names to the list, the **-a** keyletter is used. For example:

```
admin -axyz -awql s.abc
```

adds the login names “xyz” and “wql” to the list. The **-a** keyletter may be used whether *admin* is creating a new SCCS file or processing an existing one, and may appear several times. The **-e** keyletter is used in an analogous manner if one wishes to remove (“erase”) login names from the list.

5.4 prt

Prt is used to format and print on the standard output all or parts of an SCCS file (see Section 6.2), preceded by the file’s name. The portions of the file to be printed are selected by specifying certain keyletters, which, together with the output formats they generate, are fully described in *prt(1)*. This section only describes briefly the **-d**, **-u**, **-f**, and **-t** keyletters, which are sufficient to print all of the more interesting portions of an SCCS file.

The **-d** keyletter is used to print the *delta table* of an SCCS file. The delta table is that portion of the file that contains information relevant to the creation of each delta of the file, namely the *sid* of the delta, the date and time of creation, the *login* name of the creator, and the numbers of lines inserted, deleted, and unchanged by the delta. The commentary that is entered when a delta is created is also part of the delta table. Thus, executing the command:

```
prt -d s.abc
```

provides a history of the evolution of the SCCS file. In the absence of *any* keyletters, the **-d** keyletter is assumed.

The **-u** keyletter is used to print the *user list*. The **-f** keyletter causes the printing of all the *flags* of the SCCS file. The **-t** keyletter is used to print the *descriptive text* of the SCCS file (see Section 6.2); this could be used, for example, to generate a complete set of file summaries, by executing:

```
prt -t sees
```

in which “sees” is the name of a directory containing the SCCS files.

Although *prt* makes the examination of SCCS files convenient, other PWB/UNIX commands (e.g., *cd(D)*, *grep(D)*) can be used to create customized print commands in the form of Shell procedures.

5.5 help

The *help* command prints explanations of SCCS commands and of messages that these commands may print. Arguments to *help*, zero or more of which may be supplied, are simply the names of SCCS commands or the code numbers that appear in parentheses after SCCS messages. If no argument is given, *help* prompts for one. *Help* has no concept of *keyletter* arguments or *file* arguments. Explanatory information related to an argument, if it exists, is printed on the standard output. If no information is found, an error message is printed. Note that each argument is processed independently, and an error resulting from one argument will *not* terminate the processing of the other arguments.

Explanatory information related to a command is a synopsis of the command. For example:

```
help ge5 rmdel
```

produces:

```
ge5:
"nonexistent sid"
The specified sid does not exist in the
given file.
Check for typos.
```

```
rmdel:
rmdel -rSID name ...
```

5.6 rmdel

The *rmdel* command is provided to allow *removal* of a delta from an SCCS file, though its use should be reserved for those cases in which incorrect, global changes were made a part of the delta to be removed.

The delta to be removed must be a "leaf" delta. That is, it must be the latest (most recently created) delta on its branch or on the trunk of the SCCS file tree. In Figure 3, only deltas 1.3.1.2, 1.3.2.2, and 2.2 can be removed; once they are removed, then deltas 1.3.2.1 and 2.1 can be removed, and so on.

To be allowed to remove a delta, the effective user must have write permission in the directory containing the SCCS file. In addition, the real user must either be the one who created the delta being removed, or be the owner of the SCCS file and its directory.

The *-r* keyletter, which is mandatory, is used to specify the *complete* SID of the delta to be removed (i.e., it must have two components for a trunk delta, and four components for a branch delta). Thus:

```
rmdel -r2.3 s.abc
```

specifies the removal of (trunk) delta "2.3" of the SCCS file. Before removal of the delta, *rmdel* checks that the *release* number (R) of the given SID satisfies the relation:

$$\text{floor} \leq R \leq \text{ceiling}$$

In addition, the login name of the user must appear in the file's *user list*, or the *user list* must be empty. If these conditions are not satisfied, processing is terminated, and the delta is not removed. After the specified delta has been removed, its type indicator in the *delta table* of the SCCS file (see Section 6.2) is changed from "D" (for "delta") to "R" (for "removed").

5.7 chghist

The *chghist* command is used to *change* a delta's commentary that was supplied when that delta was created. Its invocation is analogous to that of the *rmdel* command, except that the delta to be processed is *not* required to be a leaf delta. For example:

```
chghist -r3.4 s.abc
```

specifies that the commentary of delta "3.4" of the SCCS file is to be changed.

The *new* commentary is solicited by *chghost* in the manner of the *delta* command. The old commentary associated with the specified delta is kept, but it is preceded by a comment line indicating that it has been changed (i.e., superseded), and the new commentary is entered ahead of this comment line. The "inserted" comment line records the login name of the user executing *chghost* and the time of its execution.

5.8 what

The *what* command is used to find identifying information within *any* PWB/UNIX file whose name is given as an argument to *what*. Directory names and a name of "-" (a lone minus sign) are *not* treated specially, as they are by other SCCS commands, and no *keyletters* are accepted by the command.

What searches the given file(s) for all occurrences of the string "@(#)", which is the replacement for the %Z% ID keyword (see *get*(1)), and prints (on the standard output) what follows that string until the first double quote ("), greater than (>), newline, or (non-printing) NUL character. Thus, for example, if the SCCS file "s.prog.c" (which is a C program), contains the following line (the %M% and %I% ID keywords were defined in Section 5.1.1):

```
char id[] "%Z%M%: %I%";
```

and then the command:

```
get -r3.4 s.prog.c
```

is executed, and finally the resulting *g-file* is compiled to produce "prog.o" and "a.out", then the command:

```
what prog.c prog.o a.out
```

produces:

```
prog.c:
  prog.c:3.4
prog.o:
  prog.c:3.4
a.out:
  prog.c:3.4
```

The string searched for by *what* need not be inserted via an ID keyword of *get*; it may be inserted in any convenient manner.

5.9 sccsdiff

The *sccsdiff* command determines (and prints on the standard output) the differences between two specified versions of one or more SCCS files. The versions to be compared are specified by using the -r keyletter, whose format is the same as for the *get* command. The two versions *must* be specified as the first two arguments to this command in the order in which they were created, i.e., the older version is specified first. Any following keyletters are interpreted as arguments to the *pr*(1) command (which actually prints the differences) and must appear before any file names. SCCS files to be processed are named last. Directory names and a name of "-" (a lone minus sign) are *not* acceptable to *sccsdiff*.

The differences are printed in the form generated by *diff*(1). The following is an example of the invocation of *sccsdiff*:

```
sccsdiff -r3.4 -r5.6 s.abc
```

5.10 comb

Comb generates a *Shell procedure* (see *sh*(1)) which attempts to reconstruct the named SCCS files so that the reconstructed files are smaller than the originals. The generated Shell procedure is written on the standard output.

Named SCCS files are reconstructed by discarding unwanted deltas and combining specified other deltas. The intended use is for those SCCS files that contain deltas that are so old that they are no longer useful. It is *not* recommended that *comb* be used as a matter of routine; its use should be restricted to a *very* small number of times in the life of an SCCS file.

In the absence of any keyletters, *comb* preserves only leaf deltas and the minimum number of ancestor deltas necessary to preserve the "shape" of the SCCS file tree. The effect of this is to eliminate "middle" deltas on the trunk and on all branches of the tree. Thus, in Figure 3, deltas 1.2, 1.3.2.1, 1.4, and 2.1 would be eliminated. Some of the keyletters are summarized as follows:

The **-p** keyletter specifies the oldest delta that is to be preserved in the reconstruction. All older deltas are discarded.

The **-e** keyletter specifies a *list* (see *get(1)* for the syntax of such a list) of deltas to be preserved. All other deltas are discarded.

The **-s** keyletter causes the generation of a Shell procedure, which, when run, produces *only* a report summarizing the percentage space (if any) to be saved by reconstructing each named SCCS file. It is recommended that *comb* be run with this keyletter (in addition to any others desired) *before* any actual reconstructions.

It should be noted that the Shell procedure generated by *comb* is *not* guaranteed to save any space. In fact, it is possible for the reconstructed file to be *larger* than the original. Note, too, that the shape of the SCCS file tree may be altered by the reconstruction process.

6. SCCS FILES

This section discusses several topics that must be considered before extensive use is made of SCCS. These topics deal with the protection mechanisms relied upon by SCCS, the format of SCCS files, and the recommended procedures for auditing SCCS files.

6.1 Protection

SCCS relies on the capabilities of the PWB/UNIX operating system for most of the protection mechanisms required to prevent unauthorized changes to SCCS files (i.e., changes made by non-SCCS commands). The only protection features provided directly by SCCS are the *release floor* and *ceiling* flags, and the *user list* (see Section 5.1.3).

New SCCS files created by the *admin* command are given mode 444 (read only). It is recommended that this mode *not* be changed, as it prevents any direct modification of the files by non-SCCS commands. It is further recommended that the directories containing SCCS files be given mode 755, which allows only the *owner* of the directory to modify its contents.

SCCS files should be kept in directories that contain only SCCS files and any temporary files created by SCCS commands. This simplifies protection and auditing of SCCS files (see Section 6.3). The contents of directories should correspond to convenient logical groupings, e.g., sub-systems of a large project.

SCCS files must have only *one* link (name). The reason for this is that those commands that modify SCCS files do so by creating a temporary copy of the file (called the *x-file*, see Section 4) and, upon completion of processing, remove the old file and rename the *x-file*. If the old file has more than one link, removing it and renaming the *x-file* would break the link. Rather than process such files, SCCS commands produce an error message. All SCCS files *must* have names that begin with "s."

When only one user (or a group of users who share the same PWB/UNIX user identification number—user ID—see *passwd(1)*) uses SCCS, the real and effective user IDs are the same, and that user ID owns the directories containing SCCS files. In addition, when several users share the same user ID (even though they may have different *login* names), all such users have identical file permissions. Therefore, SCCS may be used directly by any one of these users, without any preliminary preparation.

However, there are situations (for example, in large software development projects) in which it is not practical to give the same user ID to all users of SCCS. In these cases, one user (equivalently, one user ID) must be chosen as the "owner" of the SCCS files and be the one who will "administer" them (e.g., by using the *admin* command). This user is termed the *SCCS administrator* for that project. Because other users of SCCS do not have the same privileges and permissions as the SCCS administrator, they are not able to execute directly those commands that require write permission in the directory containing the SCCS files. Therefore, a project-dependent program is required to provide an interface to the *get*, *delta*, and, if desired, *rmdel* and *chghist* commands.

The interface program must be owned by the SCCS administrator, and must have the *set user ID on execution* bit on (see *chmod(1)*), so that the effective user ID is the user ID of the administrator. This program's function is to invoke the desired SCCS command and to cause it to *inherit* the privileges of the interface program for the duration of that command's execution. In this manner, the owner of an SCCS file can modify it at will. Other users whose *login* names are in the *user list* for that file (but who are *not* its owners) are given the necessary permissions only for the duration of the execution of the interface program, and are thus able to modify the SCCS files only through the use of *delta* and, possibly, *rmdel* and *chghist*. The project-dependent interface program, as its name implies, must be custom-built for each project.

6.2 Format

SCCS files are composed of lines of ASCII text⁷ arranged in six parts, as follows:

Checksum	A line containing the "logical" sum of all the characters of the file (<i>not</i> including this checksum itself).
Delta Table	Information about each delta, such as its type, its SID, date and time of creation, and commentary.
User Names	List of login names of users who are allowed to modify the file by adding or removing deltas.
Flags	Indicators that control certain actions of various SCCS commands.
Descriptive Text	Arbitrary text provided by the user; usually a summary of the contents and purpose of the file.
Body	Actual text that is being administered by SCCS, intermixed with internal SCCS control lines.

Detailed information about the contents of the various sections of the file may be found in *sccsfile(5)*; the *checksum* is the only portion of the file which is of interest below.

It is important to note that because SCCS files are ASCII files, they may be processed by various PWB UNIX commands, such as *ed(1)*, *grep(1)*, and *cat(1)*. This is very convenient in those instances in which an SCCS file must be modified manually (e.g., when the time and date of a delta was recorded incorrectly because the system clock was set incorrectly), or when it is desired to simply "look" at the file.

☛ *Extreme care should be exercised when modifying SCCS files with non-SCCS commands.*

6.3 Auditing

On rare occasions, perhaps due to an operating system or hardware malfunction, an SCCS file, or portions of it (i.e., one or more "blocks") can be destroyed. SCCS commands (like most PWB UNIX commands) issue an error message when a file does not exist. In addition, SCCS commands use the *checksum* stored in the SCCS file to determine whether a file has been *corrupted* since it was last accessed (possibly by having lost one or more blocks, or by having been modified with, for example, *ed(1)*). No SCCS command will process a corrupted SCCS file except the *admin* command with the **-h** or **-r** keyletters, as described below.

It is recommended that SCCS files be audited (checked) for possible corruptions on a regular basis. The simplest and fastest way to perform an audit is to execute the *admin* command with the **-h** keyletter on all SCCS files:

```
admin -h s.file1 s.file2 ...
or
admin -h directory1 directory2 ...
```

⁷ Versions of SCCS up to and including Version 3 used non-ASCII files. Therefore, files created by earlier versions of SCCS are not compatible with Version 4 of SCCS.

If the new checksum of any file is not equal to the checksum in the first line of that file, the message:

```
corrupted file (co6)
```

is produced for that file. This process continues until all the files have been examined. When examining directories (as in the second example above), the process just described will not detect *missing* files. A simple way to detect whether *any* files are missing from a directory is to periodically execute the *ls(1)* command on that directory, and compare the outputs of the most current and the previous executions. Any file whose name appears in the previous output but not in the current one has been removed by some means.

Whenever a file has been corrupted, the manner in which the file is restored depends upon the extent of the corruption. If damage is extensive, the best solution is to contact the local PWB/UNIX operations group to request a restoral of the file from a backup copy. In the case of minor damage, repair through use of the editor *ed(1)* may be possible. In the latter case, after such repair, the following command must be executed:

```
admin -z s.file
```

The purpose of this is to recompute the checksum to bring it into agreement with the actual contents of the file. After this command is executed on a file, any corruption which may have existed in that file will no longer be detectable.

REFERENCES

- [1] Ritchie, D. M., and Thompson, K. The UNIX Time-Sharing System. *Comm. ACM* 17(7):365-75, July 1974.
- [2] Kernighan, B. W. UNIX for Beginners. Bell Laboratories, 1973.
- [3] Kernighan, B. W. A Tutorial Introduction to the UNIX Text Editor. Bell Laboratories, 1973.
- [4] Dolotta, T. A., Haight, R. C., and Piskorik, E. M., eds. *PWB/UNIX User's Manual—Edition 1.0*. Bell Laboratories, May 1977.
- [5] Kernighan, B. W., and Ritchie, D. M. UNIX Programming. Bell Laboratories, 1973.

NROFF/TROFF User's Manual

Joseph F. Ossanna

Bell Laboratories
Murray Hill, New Jersey 07974

Introduction

NROFF and TROFF are text processors under the PDP-11 UNIX Time-Sharing System¹ that format text for typewriter-like terminals and for a Graphic Systems phototypesetter, respectively. They accept lines of text interspersed with lines of format control information and format the text into a printable, paginated document having a user-designed style. NROFF and TROFF offer unusual freedom in document styling, including: arbitrary style headers and footers; arbitrary style footnotes; multiple automatic sequence numbering for paragraphs, sections, etc; multiple column output; dynamic font and point-size control; arbitrary horizontal and vertical local motions at any point; and a family of automatic overstriking, bracket construction, and line drawing functions.

NROFF and TROFF are highly compatible with each other and it is almost always possible to prepare input acceptable to both. Conditional input is provided that enables the user to embed input expressly destined for either program. NROFF can prepare output directly for a variety of terminal types and is capable of utilizing the full resolution of each terminal.

Usage

The general form of invoking NROFF (or TROFF) at UNIX command level is

nroff *options files* (or **troff** *options files*)

where *options* represents any of a number of option arguments and *files* represents the list of files containing the document to be formatted. An argument consisting of a single minus (−) is taken to be a file name corresponding to the standard input. If no file names are given input is taken from the standard input. The options, which may appear in any order so long as they appear before the files, are:

Option	Effect
−olist	Print only pages whose page numbers appear in <i>list</i> , which consists of comma-separated numbers and number ranges. A number range has the form $N-M$ and means pages N through M ; a initial $-N$ means from the beginning to page N ; and a final $N-$ means from N to the end.
−nN	Number first generated page N .
−sN	Stop every N pages. NROFF will halt prior to every N pages (default $N=1$) to allow paper loading or changing, and will resume upon receipt of a newline. TROFF will stop the phototypesetter every N pages, produce a trailer to allow changing cassettes, and will resume after the phototypesetter START button is pressed.
−mname	Prepends the macro file <code>/usr/lib/tmac.name</code> to the input <i>files</i> .
−raN	Register a (one-character) is set to N .
−i	Read standard input after the input files are exhausted.
−q	Invoke the simultaneous input-output mode of the rd request.

NROFF Only

- Tname** Specifies the name of the output terminal type. Currently defined names are **37** for the (default) Model 37 teletype, **tn300** for the GE TermiNet 300 (or any terminal without half-line capabilities), **300S** for the DASI-300S, **300** for the DASI-300, and **450** for the DASI-450 (Diablo Hyterm).
- e** Produce equally-spaced words in adjusted lines, using full terminal resolution.

TROFF Only

- t** Direct output to the standard output instead of the phototypesetter.
- f** Refrain from feeding out paper and stopping phototypesetter at the end of the run.
- w** Wait until phototypesetter is available, if currently busy.
- b** TROFF will report whether the phototypesetter is busy or available. No text processing is done.
- a** Send a printable (ASCII) approximation of the results to the standard output.
- pN** Print all characters in point size *N* while retaining all prescribed spacings and motions, to reduce phototypesetter elapsed time.
- g** Prepare output for the Murray Hill Computation Center phototypesetter and direct it to the standard output.

Each option is invoked as a separate argument; for example,

```
nroff -o4,8-10 -T300S -mabc file1 file2
```

requests formatting of pages 4, 8, 9, and 10 of a document contained in the files named *file1* and *file2*, specifies the output terminal as a DASI-300S, and invokes the macro package *abc*.

Various pre- and post-processors are available for use with NROFF and TROFF. These include the equation preprocessors NEQN and EQN² (for NROFF and TROFF respectively), and the table-construction preprocessor TBL³. A reverse-line postprocessor COL⁴ is available for multiple-column NROFF output on terminals without reverse-line ability; COL expects the Model 37 Teletype escape sequences that NROFF produces by default. TK⁴ is a 37 Teletype simulator postprocessor for printing NROFF output on a Tektronix 4014. TCAT⁴ is phototypesetter-simulator postprocessor for TROFF that produces an approximation of phototypesetter output on a Tektronix 4014. For example, in

```
tbl files | eqn | troff -t options | tcat
```

the first | indicates the piping of TBL's output to EQN's input; the second the piping of EQN's output to TROFF's input; and the third indicates the piping of TROFF's output to TCAT. GCAT⁴ can be used to send TROFF (-g) output to the Murray Hill Computation Center.

The remainder of this manual consists of: a Summary and Index; a Reference Manual keyed to the index; and a set of Tutorial Examples. Another tutorial is [5].

Joseph F. Ossanna

References

- [1] K. Thompson, D. M. Ritchie, *UNIX Programmer's Manual*, Sixth Edition (May 1975).
- [2] B. W. Kernighan, L. L. Cherry, *Typesetting Mathematics — User's Guide (Second Edition)*, Bell Laboratories internal memorandum.
- [3] M. E. Lesk, *Tbl — A Program to Format Tables*, Bell Laboratories internal memorandum.
- [4] Internal on-line documentation, on UNIX.
- [5] B. W. Kernighan, *A TROFF Tutorial*, Bell Laboratories internal memorandum.

SUMMARY AND INDEX

Request Form	Initial Value*	If No Argument	Notes#	Explanation
1. General Explanation				
2. Font and Character Size Control				
.ps ± N	10 point	previous	E	Point size; also \s±N.†
.ss N	12/36 em	ignored	E	Space-character size set to N/36 em.†
.cs FNM	off	-	P	Constant character space (width) mode (font F).†
.bd FN	off	-	P	Embolden font F by N-1 units.†
.bd S FN	off	-	P	Embolden Special Font when current font is F.†
.ft F	Roman	previous	E	Change to font F = x, xx, or 1-4. Also \fx, \f(xx, \fN.
.fp NF	R,I,B,S	ignored	-	Font named F mounted on physical position 1 ≤ N ≤ 4.
3. Page Control				
.pl ± N	11 in	11 in	v	Page length.
.bp ± N	N=1	-	B‡,v	Eject current page; next page number N.
.pn ± N	N=1	ignored	-	Next page number N.
.po ± N	0; 26/27 in	previous	v	Page offset.
.ne N	-	N=1 V	D,v	Need N vertical space (V = vertical spacing).
.mk R	none	internal	D	Mark current vertical place in register R.
.rt ± N	none	internal	D,v	Return (<i>upward only</i>) to marked vertical place.
4. Text Filling, Adjusting, and Centering				
.br	-	-	B	Break.
.fi	fill	-	B,E	Fill output lines.
.nf	fill	-	B,E	No filling or adjusting of output lines.
.ad c	adj,both	adjust	E	Adjust output lines with mode c.
.na	adjust	-	E	No output line adjusting.
.ce N	off	N=1	B,E	Center following N input text lines.
5. Vertical Spacing				
.vs N	1/6in;12pts	previous	E,p	Vertical base line spacing (V).
.ls N	N=1	previous	E	Output N-1 Vs after each text output line.
.sp N	-	N=1 V	B,v	Space vertical distance N in either direction.
.sv N	-	N=1 V	v	Save vertical distance N.
.os	-	-	-	Output saved vertical distance.
.ns	space	-	D	Turn no-space mode on.
.rs	-	-	D	Restore spacing; turn no-space mode off.
6. Line Length and Indenting				
.ll ± N	6.5 in	previous	E,m	Line length.
.in ± N	N=0	previous	B,E,m	Indent.
.ti ± N	-	ignored	B,E,m	Temporary indent.
7. Macros, Strings, Diversion, and Position Traps				
.de xx yy	-	.yy=..	-	Define or redefine macro xx; end at call of yy.
.am xx yy	-	.yy=..	-	Append to a macro.
.ds xx string	-	ignored	-	Define a string xx containing string.
.as xx string	-	ignored	-	Append string to string xx.

*Values separated by ";" are for NROFF and TROFF respectively.

#Notes are explained at the end of this Summary and Index

†No effect in NROFF.

‡The use of " " as control character (instead of ".") suppresses the break function.

Request Form	Initial Value	If No Argument	Notes	Explanation
.rm <i>xx</i>	-	ignored	-	Remove request, macro, or string.
.rn <i>xx yy</i>	-	ignored	-	Rename request, macro, or string <i>xx</i> to <i>yy</i> .
.di <i>xx</i>	-	end	D	Divert output to macro <i>xx</i> .
.da <i>xx</i>	-	end	D	Divert and append to <i>xx</i> .
.wh <i>N xx</i>	-	-	v	Set location trap; negative is w.r.t. page bottom.
.ch <i>xx N</i>	-	-	v	Change trap location.
.dt <i>N xx</i>	-	off	D,v	Set a diversion trap.
.it <i>N xx</i>	-	off	E	Set an input-line count trap.
.em <i>xx</i>	none	none	-	End macro is <i>xx</i> .

8. Number Registers

.nr <i>R ± N M</i>	-	-	u	Define and set number register <i>R</i> ; auto-increment by <i>M</i> .
.af <i>R c</i>	arabic	-	-	Assign format to register <i>R</i> (<i>c</i> =1, i, I, a, A).
.rr <i>R</i>	-	-	-	Remove register <i>R</i> .

9. Tabs, Leaders, and Fields

.ta <i>Nt ...</i>	0.8; 0.5in	none	E,m	Tab settings; <i>left</i> type, unless <i>t</i> =R(right), C(centered).
.tc <i>c</i>	none	none	E	Tab repetition character.
.lc <i>c</i>	.	none	E	Leader repetition character.
.fc <i>a b</i>	off	off	-	Set field delimiter <i>a</i> and pad character <i>b</i>

10. Input and Output Conventions and Character Translations

.ec <i>c</i>	\	\	-	Set escape character.
.eo	on	-	-	Turn off escape character mechanism.
.lg <i>N</i>	-; on	on	-	Ligature mode on if <i>N</i> >0.
.ul <i>N</i>	off	<i>N</i> =1	E	Underline (italicize in TROFF) <i>N</i> input lines.
.cu <i>N</i>	off	<i>N</i> =1	E	Continuous underline in NROFF; like ul in TROFF.
.uf <i>F</i>	Italic	Italic	-	Underline font set to <i>F</i> (to be switched to by ul).
.cc <i>c</i>	:	:	E	Set control character to <i>c</i> .
.c2 <i>c</i>	:	:	E	Set nobreak control character to <i>c</i> .
.tr <i>abcd... </i>	none	-	O	Translate <i>a</i> to <i>b</i> , etc. on output.

11. Local Horizontal and Vertical Motions, and the Width Function

12. Overstrike, Bracket, Line-drawing, and Zero-width Functions

13. Hyphenation.

.nh	hyphenate	-	E	No hyphenation.
.hy <i>N</i>	hyphenate	hyphenate	E	Hyphenate; <i>N</i> = mode.
.hc <i>c</i>	\%	\%	E	Hyphenation indicator character <i>c</i> .
.hw <i>wordl ...</i>		ignored	-	Exception words.

14. Three Part Titles.

.tl ' <i>left center right</i> '	-	-	-	Three part title.
.pc <i>c</i>	%	off	-	Page number character.
.lt $\pm N$	6.5in	previous	E,m	Length of title.

15. Output Line Numbering.

.nm $\pm N M S I$	off	-	E	Number mode on or off, set parameters.
.nn <i>N</i>	-	<i>N</i> =1	E	Do not number next <i>N</i> lines.

16. Conditional Acceptance of Input

.if <i>c anything</i>	-	-	-	If condition <i>c</i> true, accept <i>anything</i> as input, for multi-line use $\{anything\}$.
------------------------------	---	---	---	--

Request Form	Initial Value	If No Argument	Notes	Explanation
.if !c anything	-	-	-	If condition <i>c</i> false, accept <i>anything</i> .
.if N anything	-	-	u	If expression $N > 0$, accept <i>anything</i> .
.if !N anything	-	-	u	If expression $N \leq 0$, accept <i>anything</i> .
.if 'string1' string2' anything	-	-	-	If <i>string1</i> identical to <i>string2</i> , accept <i>anything</i> .
.if !'string1' string2' anything	-	-	-	If <i>string1</i> not identical to <i>string2</i> , accept <i>anything</i> .
.ie c anything	-	-	u	If portion of if-else; all above forms (like if).
.el anything	-	-	-	Else portion of if-else.
17. Environment Switching.				
.ev N	<i>N=0</i>	previous	-	Environment switched (<i>push down</i>).
18. Insertions from the Standard Input				
.rd prompt	-	<i>prompt=BEL</i>	-	Read insertion.
.ex	-	-	-	Exit from NROFF/TROFF.
19. Input/Output File Switching				
.so filename	-	-	-	Switch source file (<i>push down</i>).
.nx filename	-	end-of-file	-	Next file.
.pi program	-	-	-	Pipe output to <i>program</i> (NROFF only).
20. Miscellaneous				
.mc c N	-	off	E,m	Set margin character <i>c</i> and separation <i>N</i> .
.tm string	-	newline	-	Print <i>string</i> on terminal (UNIX standard message output).
.ig yy	-	<i>.yy=..</i>	-	Ignore till call of <i>yy</i> .
.pm t	-	all	-	Print macro names and sizes; if <i>t</i> present, print only total of sizes.
.fl	-	-	B	Flush output buffer.
21. Output and Error Messages				

Notes-

- B** Request normally causes a break.
- D** Mode or relevant parameters associated with current diversion level.
- E** Relevant parameters are a part of the current environment.
- O** Must stay in effect until logical output.
- P** Mode must be still or again in effect at the time of physical output.
- v,p,m,u** Default scale indicator; if not specified, scale indicators are *ignored*.

Alphabetical Request and Section Number Cross Reference

ad 4	cc 10	ds 7	fc 9	ie 16	ll 6	nh 13	pi 19	rn 7	ta 9	vs 5
af 8	ce 4	dt 7	fi 4	if 16	ls 5	nm 15	pl 3	rr 8	tc 9	wh 7
am 7	ch 7	ec 10	fl 20	ig 20	lt 14	nn 15	pm 20	rs 5	ti 6	
as 7	cs 2	el 16	fp 2	in 6	mc 20	nr 8	pn 3	rt 3	tl 14	
bd 2	cu 10	em 7	ft 2	it 7	mk 3	ns 5	po 3	so 19	tm 20	
bp 3	da 7	eo 10	hc 13	lc 9	na 4	nx 19	ps 2	sp 5	tr 10	
br 4	de 7	ev 17	hw 13	lg 10	ne 3	os 5	rd 18	ss 2	uf 10	
c2 10	di 7	ex 18	hy 13	li 10	nf 4	pc 14	rm 7	sv 5	ul 10	

Escape Sequences for Characters, Indicators, and Functions

Section Reference	Escape Sequence	Meaning
10.1	\\	\ (to prevent or delay the interpretation of \)
10.1	\e	Printable version of the <i>current</i> escape character.
2.1	\`	` (acute accent); equivalent to \(\aa
2.1	\`	` (grave accent); equivalent to \(\ga
2.1	\-	- Minus sign in the <i>current</i> font
7	\.	Period (dot) (see de)
11.1	\(space)	Unpaddable space-size space character
11.1	\0	Digit width space
11.1	\	1/6 em narrow space character (zero width in NROFF)
11.1	\^	1/12 em half-narrow space character (zero width in NROFF)
4.1	\&	Non-printing, zero width character
10.6	\!	Transparent line indicator
10.7	\"	Beginning of comment
7.3	\\$N	Interpolate argument $1 \leq N \leq 9$
13	\%	Default optional hyphenation character
2.1	\(xx	Character named <i>xx</i>
7.1	\.x, *(xx	Interpolate string <i>x</i> or <i>xx</i>
9.1	\a	Non-interpreted leader character
12.3	\b'abc...'	Bracket building function
4.2	\c	Interrupt text processing
11.1	\d	Forward (down) 1/2 em vertical motion (1/2 line in NROFF)
2.2	\fx,\f(xx,\fN	Change to font named <i>x</i> or <i>xx</i> , or position <i>N</i>
11.1	\h'N'	Local horizontal motion; move right <i>N</i> (<i>negative left</i>)
11.3	\kx	Mark horizontal <i>input</i> place in register <i>x</i>
12.4	\l'Nc'	Horizontal line drawing function (optionally with <i>c</i>)
12.4	\L'Nc'	Vertical line drawing function (optionally with <i>c</i>)
8	\nx,\n(xx	Interpolate number register <i>x</i> or <i>xx</i>
12.1	\o'abc...'	Overstrike characters <i>a, b, c, ...</i>
4.1	\p	Break and spread output line
11.1	\r	Reverse 1 em vertical motion (reverse line in NROFF)
2.3	\sN,\s±N	Point-size change function
9.1	\t	Non-interpreted horizontal tab
11.1	\u	Reverse (up) 1/2 em vertical motion (1/2 line in NROFF)
11.1	\v'N'	Local vertical motion; move down <i>N</i> (<i>negative up</i>)
11.2	\w'string'	Interpolate width of <i>string</i>
5.2	\x'N'	Extra line-space function (<i>negative before, positive after</i>)
12.2	\zc	Print <i>c</i> with zero width (without spacing)
16	\{	Begin conditional input
16	\}	End conditional input
10.7	\(newline)	Concealed (ignored) newline
	\X	<i>X</i> , any character <i>not</i> listed above

The escape sequences \\, \., \", \\$, *, \a, \n, \t, and \(\newline) are interpreted in *copy mode* (§7.2).

Predefined General Number Registers

<i>Section Reference</i>	<i>Register Name</i>	<i>Description</i>
3	%	Current page number.
11.2	ct	Character type (set by <i>width</i> function).
7.4	dl	Width (maximum) of last completed diversion.
7.4	dn	Height (vertical size) of last completed diversion.
-	dw	Current day of the week (1-7).
-	dy	Current day of the month (1-31).
11.3	hp	Current horizontal place on <i>input</i> line.
15	ln	Output line number.
-	mo	Current month (1-12).
4.1	nl	Vertical position of last printed text base-line.
11.2	sb	Depth of string below base line (generated by <i>width</i> function).
11.2	st	Height of string above base line (generated by <i>width</i> function).
-	yr	Last two digits of current year.

Predefined Read-Only Number Registers

<i>Section Reference</i>	<i>Register Name</i>	<i>Description</i>
7.3	.S	Number of arguments available at the current macro level.
-	.A	Set to 1 in TROFF, if <i>-a</i> option used; always 1 in NROFF.
11.1	.H	Available horizontal resolution in basic units.
-	.T	Set to 1 in NROFF, if <i>-T</i> option used; always 0 in TROFF.
11.1	.V	Available vertical resolution in basic units.
5.2	.a	Post-line extra line-space most recently utilized using <i>\x'N'</i> .
-	.c	Number of <i>lines</i> read from current input file.
7.4	.d	Current vertical place in current diversion; equal to <i>nl</i> , if no diversion.
2.2	.f	Current font as physical quadrant (1-4).
4	.h	Text base-line high-water mark on current page or diversion.
6	.i	Current indent.
6	.l	Current line length.
4	.n	Length of text portion on previous output line.
3	.o	Current page offset.
3	.p	Current page length.
2.3	.s	Current point size.
7.5	.t	Distance to the next trap.
4.1	.u	Equal to 1 in fill mode and 0 in nofill mode.
5.1	.v	Current vertical line spacing.
11.2	.w	Width of previous character.
-	.x	Reserved version-dependent register.
-	.y	Reserved version-dependent register.
7.4	.z	Name of current diversion.

REFERENCE MANUAL

1. General Explanation

1.1. Form of input. Input consists of *text lines*, which are destined to be printed, interspersed with *control lines*, which set parameters or otherwise control subsequent processing. Control lines begin with a *control character*—normally . (period) or ` (acute accent)—followed by a one or two character name that specifies a basic *request* or the substitution of a user-defined *macro* in place of the control line. The control character ` suppresses the *break* function—the forced output of a partially filled line—caused by certain requests. The control character may be separated from the request/macro name by white space (spaces and/or tabs) for esthetic reasons. Names must be followed by either space or newline. Control lines with unrecognized names are ignored.

Various special functions may be introduced anywhere in the input by means of an *escape* character, normally \. For example, the function \nR causes the interpolation of the contents of the *number register* R in place of the function; here R is either a single character name as in \n.x, or left-parenthesis-introduced, two-character name as in \n(xx).

1.2. Formatter and device resolution. TROFF internally uses 432 units/inch, corresponding to the Graphic Systems phototypesetter which has a horizontal resolution of 1/432 inch and a vertical resolution of 1/144 inch. NROFF internally uses 240 units/inch, corresponding to the least common multiple of the horizontal and vertical resolutions of various typewriter-like output devices. TROFF rounds horizontal/vertical numerical parameter input to the actual horizontal/vertical resolution of the Graphic Systems typesetter. NROFF similarly rounds numerical input to the actual resolution of the output device indicated by the -T option (default Model 37 Teletype).

1.3. Numerical parameter input. Both NROFF and TROFF accept numerical input with the appended scale indicators shown in the following table, where S is the current type size in points, V is the current vertical line spacing in basic units, and C is a *nominal character width* in basic units.

Scale Indicator	Meaning	Number of basic units	
		TROFF	NROFF
i	Inch	432	240
c	Centimeter	432×50/127	240×50/127
P	Pica = 1/6 inch	72	240/6
m	Em = S points	6×S	C
n	En = Em/2	3×S	C, same as Em
p	Point = 1/72 inch	6	240/72
u	Basic unit	1	1
v	Vertical line space	V	V
none	Default, see below		

In NROFF, both the em and the en are taken to be equal to the C, which is output-device dependent; common values are 1/10 and 1/12 inch. Actual character widths in NROFF need not be all the same and constructed characters such as -> (→) are often extra wide. The default scaling is ems for the horizontally-oriented requests and functions ll, in, ti, ta, lt, po, mc, \h, and \l; Vs for the vertically-oriented requests and functions pl, wh, ch, dt, sp, sv, ne, rt, \v, \x, and \L; p for the vs request; and u for the requests nr, if, and ie. All other requests ignore any scale indicators. When a number register containing an already appropriately scaled number is interpolated to provide numerical input, the unit scale indicator u may need to be appended to prevent an additional inappropriate default scaling.

The number, N , may be specified in decimal-fraction form but the parameter finally stored is rounded to an integer number of basic units.

The *absolute position* indicator | may be prepended to a number N to generate the distance to the vertical or horizontal place N . For vertically-oriented requests and functions, | N becomes the distance in basic units from the current vertical place on the page or in a *diversion* (§7.4) to the vertical place N . For all other requests and functions, | N becomes the distance from the current horizontal place on the *input* line to the horizontal place N . For example,

`.sp |3.2c`

will space *in the required direction* to 3.2 centimeters from the top of the page.

1.4. Numerical expressions. Wherever numerical input is expected an expression involving parentheses, the arithmetic operators +, -, /, *, % (mod), and the logical operators <, >, <=, >=, = (or ==), & (and), : (or) may be used. Except where controlled by parentheses, evaluation of expressions is left-to-right; there is no operator precedence. In the case of certain requests, an initial + or - is stripped and interpreted as an increment or decrement indicator respectively. In the presence of default scaling, the desired scale indicator must be attached to every number in an expression for which the desired and default scaling differ. For example, if the number register x contains 2 and the current point size is 10, then

`.ll (4.25i+\nxP+3)/2u`

will set the line length to 1/2 the sum of 4.25 inches + 2 picas + 30 points.

1.5. Notation. Numerical parameters are indicated in this manual in two ways. $\pm N$ means that the argument may take the forms N , $+N$, or $-N$ and that the corresponding effect is to set the affected parameter to N , to increment it by N , or to decrement it by N respectively. Plain N means that an initial algebraic sign is *not* an increment indicator, but merely the sign of N . Generally, unreasonable numerical input is either ignored or truncated to a reasonable value. For example, most requests expect to set parameters to non-negative values; exceptions are **sp**, **wh**, **ch**, **nr**, and **if**. The requests **ps**, **ft**, **po**, **vs**, **ls**, **ll**, **in**, and **lt** restore the *previous* parameter value in the *absence* of an argument.

Single character arguments are indicated by single lower case letters and one/two character arguments are indicated by a pair of lower case letters. Character string arguments are indicated by multi-character mnemonics.

2. Font and Character Size Control

2.1. Character set. The TROFF character set consists of the Graphics Systems Commercial II character set plus a Special Mathematical Font character set—each having 102 characters. These character sets are shown in the attached Table I. All ASCII characters are included, with some on the Special Font. With three exceptions, the ASCII characters are input as themselves, and non-ASCII characters are input in the form \ xx where xx is a two-character name given in the attached Table II. The three ASCII exceptions are mapped as follows:

ASCII Input		Printed by TROFF	
Character	Name	Character	Name
'	acute accent	'	close quote
`	grave accent	'	open quote
-	minus	-	hyphen

The characters ' , ` , and - may be input by \', \`, and \- respectively or by their names (Table II). The ASCII characters @, #, ", ' , ` , < , > , \ , { , } , ~ , ? , and _ exist only on the Special Font and are printed as a 1-em space if that Font is not mounted.

NROFF understands the entire TROFF character set, but can in general print only ASCII characters, additional characters as may be available on the output device, such characters as may be able to be constructed by overstriking or other combination, and those that can reasonably be mapped into other printable characters. The exact behavior is determined by a driving table prepared for each device. The

characters `'`, ```, and `_` print as themselves.

2.2. Fonts. The default mounted fonts are Times Roman (**R**), Times Italic (**I**), Times Bold (**B**), and the Special Mathematical Font (**S**) on physical typesetter positions 1, 2, 3, and 4 respectively. These fonts are used in this document. The *current* font, initially Roman, may be changed (among the mounted fonts) by use of the `ft` request, or by imbedding at any desired point either `\fx`, `\f(xx)`, or `\fN` where `x` and `xx` are the name of a mounted font and `N` is a numerical font position. It is *not* necessary to change to the Special font; characters on that font are automatically handled. A request for a named but not-mounted font is *ignored*. TROFF can be informed that any particular font is mounted by use of the `fp` request. The list of known fonts is installation dependent. In the subsequent discussion of font-related requests, `F` represents either a one/two-character font name or the numerical font position, 1-4. The current font is available (as numerical position) in the read-only number register `.f`.

NROFF understands font control and normally underlines Italic characters (see §10.5).

2.3. Character size. Character point sizes available on the Graphic Systems typesetter are 6, 7, 8, 9, 10, 11, 12, 14, 16, 18, 20, 22, 24, 28, and 36. This is a range of 1/12 inch to 1/2 inch. The `ps` request is used to change or restore the point size. Alternatively the point size may be changed between any two characters by imbedding a `\sN` at the desired point to set the size to `N`, or a `\s±N` ($1 \leq N \leq 9$) to increment/decrement the size by `N`; `\s0` restores the *previous* size. Requested point size values that are between two valid sizes yield the larger of the two. The current size is available in the `.s` register. NROFF ignores type size control.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes*</i>	<i>Explanation</i>
<code>.ps ±N</code>	10 point	previous	E	Point size set to $\pm N$. Alternatively imbed <code>\sN</code> or <code>\s±N</code> . Any positive size value may be requested; if invalid, the next larger valid size will result, with a maximum of 36. A paired sequence $+N, -N$ will work because the previous requested value is also remembered. Ignored in NROFF.
<code>.ss N</code>	12/36 em	ignored	E	Space-character size is set to $N/36$ ems. This size is the minimum word spacing in adjusted text. Ignored in NROFF.
<code>.cs FNM</code>	off		P	Constant character space (width) mode is set on for font <code>F</code> (if mounted); the width of every character will be taken to be $N/36$ ems. If <code>M</code> is absent, the em is that of the character's point size; if <code>M</code> is given, the em is <code>M</code> -points. All affected characters are centered in this space. Special Font characters occurring while the current font is <code>F</code> are also so treated. If <code>N</code> is absent, the mode is turned off. The mode must be still or again in effect when the characters are physically printed. Ignored in NROFF.
<code>.bd FN</code>	off		P	The characters in font <code>F</code> will be artificially emboldened by printing each one twice, separated by $N-1$ basic units. A reasonable value for <code>N</code> is 3 when the character size is in the vicinity of 10 points. If <code>N</code> is missing the embolden mode is turned off. The column heads above were printed with <code>.bd I 3</code> . The mode must be still or again in effect when the characters are physically printed. Ignored in NROFF.

*Notes are explained at the end of the Summary and Index above.

.bd <i>S F N</i>	off		P	The characters in the Special Font will be emboldened whenever the current font is <i>F</i> . This manual was printed with .bdSB3 . The mode must be still or again in effect when the characters are physically printed.
.ft <i>F</i>	Roman	previous	E	Font changed to <i>F</i> . Alternatively, imbed <code>\f<i>F</i></code> . The font name P is reserved to mean the previous font.
.fp <i>N F</i>	R,I,B,S	ignored	-	Font position. This is a statement that a font named <i>F</i> is mounted on position <i>N</i> (1-4). It is a fatal error if <i>F</i> is not known. The phototypesetter has four fonts physically mounted. Each font consists of a film strip which can be mounted on a numbered quadrant of a wheel. The default mounting sequence assumed by TROFF is R, I, B, and S on positions 1, 2, 3 and 4.

3. Page control

Top and bottom margins are *not* automatically provided; it is conventional to define two *macros* and to set *traps* for them at vertical positions 0 (top) and $-N$ (N from the bottom). See §7 and Tutorial Examples §T2. A pseudo-page transition onto the *first* page occurs either when the first *break* occurs or when the first *non-diverted* text processing occurs. Arrangements for a trap to occur at the top of the first page must be completed before this transition. In the following, references to the *current diversion* (§7.4) mean that the mechanism being described works during both ordinary and diverted output (the former considered as the top diversion level).

The useable page width on the Graphic Systems phototypesetter is about 7.54 inches, beginning about 1/27 inch from the left edge of the 8 inch wide, continuous roll paper. The physical limitations on NROFF output are output-device dependent.

Request Form	Initial Value	If No Argument	Notes	Explanation
.pl $\pm N$	11 in	11 in	v	Page length set to $\pm N$. The internal limitation is about 75 inches in TROFF and about 136 inches in NROFF. The current page length is available in the .p register.
.bp $\pm N$	$N=1$	-	B*,v	Begin page. The current page is ejected and a new page is begun. If $\pm N$ is given, the new page number will be $\pm N$. Also see request ns .
.pn $\pm N$	$N=1$	ignored	-	Page number. The next page (when it occurs) will have the page number $\pm N$. A pn must occur before the initial pseudo-page transition to effect the page number of the first page. The current page number is in the % register.
.po $\pm N$	0; 26/27 in†	previous	v	Page offset. The current <i>left margin</i> is set to $\pm N$. The TROFF initial value provides about 1 inch of paper margin including the physical typesetter margin of 1/27 inch. In TROFF the maximum (line-length) + (page-offset) is about 7.54 inches. See §6. The current page offset is available in the .o register.
.ne <i>N</i>	-	$N=1$ V	D,v	Need N vertical space. If the distance, D , to the next trap position (see §7.5) is less than N , a forward vertical space of size D occurs, which will spring the trap. If there are no remaining traps on the page, D is the

*The use of " " as control character (instead of ".") suppresses the break function.

†Values separated by ";" are for NROFF and TROFF respectively.

distance to the bottom of the page. If $D < V$, another line could still be output and spring the trap. In a diversion, D is the distance to the *diversion trap*, if any, or is very large.

<code>.mk R</code>	none	internal	D	Mark the <i>current</i> vertical place in an internal register (both associated with the current diversion level), or in register R , if given. See <code>rt</code> request.
<code>.rt ±N</code>	none	internal	D,v	Return <i>upward only</i> to a marked vertical place in the current diversion. If $±N$ (w.r.t. current place) is given, the place is $±N$ from the top of the page or diversion or, if N is absent, to a place marked by a previous <code>mk</code> . Note that the <code>sp</code> request (§5.3) may be used in all cases instead of <code>rt</code> by spacing to the absolute place stored in an explicit register; e. g. using the sequence <code>.mk Rsp \nRu</code> .

4. Text Filling, Adjusting, and Centering

4.1. Filling and adjusting. Normally, words are collected from input text lines and assembled into an output text line until some word doesn't fit. An attempt is then made to hyphenate the word in effort to assemble a part of it into the output line. The spaces between the words on the output line are then increased to spread out the line to the current *line length* minus any current *indent*. A *word* is any string of characters delimited by the *space* character or the beginning/end of the input line. Any adjacent pair of words that must be kept together (neither split across output lines nor spread apart in the adjustment process) can be tied together by separating them with the *unpaddable space* character "\ " (backslash-space). The adjusted word spacings are uniform in TROFF and the minimum interword spacing can be controlled with the `ss` request (§2). In NROFF, they are normally nonuniform because of quantization to character-size spaces; however, the command line option `-e` causes uniform spacing with full output device resolution. Filling, adjustment, and hyphenation (§13) can all be prevented or controlled. The *text length* on the last line output is available in the `.n` register, and text base-line position on the page for this line is in the `.nl` register. The text base-line high-water mark (lowest place) on the current page is in the `.h` register.

An input text line ending with `.`, `?`, or `!` is taken to be the end of a *sentence*, and an additional space character is automatically provided during filling. Multiple inter-word space characters found in the input are retained, except for trailing spaces; initial spaces also cause a *break*.

When filling is in effect, a `\p` may be imbedded or attached to a word to cause a *break* at the *end* of the word and have the resulting output line *spread out* to fill the current line length.

A text input line that happens to begin with a control character can be made to not look like a control line by prefacing it with the non-printing, zero-width filler character `\&`. Still another way is to specify output translation of some convenient character into the control character using `tr` (§10.5).

4.2. Interrupted text. The copying of an input line in *nofill* (non-fill) mode can be *interrupted* by terminating the partial line with a `\c`. The *next* encountered input text line will be considered to be a continuation of the same line of input text. Similarly, a word within *filled* text may be interrupted by terminating the word (and line) with `\c`; the next encountered text will be taken as a continuation of the interrupted word. If the intervening control lines cause a *break*, any partial line will be forced out along with any partial word.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
<code>.br</code>	-	-	E	Break. The filling of the line currently being collected is stopped and the line is output without adjustment. Text lines beginning with space characters and empty text lines (blank lines) also cause a break.

.fi	fill on	-	B,E	Fill subsequent output lines. The register .u is 1 in fill mode and 0 in nofill mode.
.nf	fill on	-	B,E	Nofill. Subsequent output lines are <i>neither filled nor adjusted</i> . Input text lines are copied directly to output lines <i>without regard</i> for the current line length.
.ad c	adj,both	adjust	E	Line adjustment is begun. If fill mode is not on, adjustment will be deferred until fill mode is back on. If the type indicator <i>c</i> is present, the adjustment type is changed as shown in the following table.

Indicator	Adjust Type
l	adjust left margin only
r	adjust right margin only
c	center
b or n	adjust both margins
absent	unchanged

.na	adjust	-	E	Noadjust. Adjustment is turned off; the right margin will be ragged. The adjustment type for ad is not changed. Output line filling still occurs if fill mode is on.
.ce N	off	N=1	B,E	Center the next <i>N</i> input text lines within the current (line-length minus indent). If <i>N</i> =0, any residual count is cleared. A break occurs after each of the <i>N</i> input lines. If the input line is too long, it will be left adjusted.

5. Vertical Spacing

5.1. Base-line spacing. The vertical spacing (*V*) between the base-lines of successive output lines can be set using the **vs** request with a resolution of 1/144 inch = 1/2 point in TROFF, and to the output device resolution in NROFF. *V* must be large enough to accommodate the character sizes on the affected output lines. For the common type sizes (9-12 points), usual typesetting practice is to set *V* to 2 points greater than the point size; TROFF default is 10-point type on a 12-point spacing (as in this document). The current *V* is available in the **.v** register. Multiple-*V* line separation (e.g. double spacing) may be requested with **ls**.

5.2. Extra line-space. If a word contains a vertically tall construct requiring the output line containing it to have extra vertical space before and/or after it, the *extra-line-space* function **\x'N'** can be imbedded in or attached to that word. In this and other functions having a pair of delimiters around their parameter (here **'**), the delimiter choice is arbitrary, except that it can't look like the continuation of a number expression for *N*. If *N* is negative, the output line containing the word will be preceded by *N* extra vertical space; if *N* is positive, the output line containing the word will be followed by *N* extra vertical space. If successive requests for extra space apply to the same line, the maximum values are used. The most recently utilized post-line extra line-space is available in the **.a** register.

5.3. Blocks of vertical space. A block of vertical space is ordinarily requested using **sp**, which honors the *no-space* mode and which does not space *past* a trap. A contiguous block of vertical space may be reserved using **sv**.

Request Form	Initial Value	If No Argument	Notes	Explanation
.vs N	1/6in;12pts	previous	E,p	Set vertical base-line spacing size <i>V</i> . Transient <i>extra</i> vertical space available with \x'N' (see above).
.ls N	N=1	previous	E	Line spacing set to $\pm N$. <i>N</i> -1 <i>V</i> s (<i>blank lines</i>) are appended to each output text line. Appended blank lines are omitted, if the text or previous appended blank line

				reached a trap position.
.sp <i>N</i>		<i>N=1 V</i>	B,v	Space vertically in <i>either</i> direction. If <i>N</i> is negative, the motion is <i>backward</i> (upward) and is limited to the distance to the top of the page. Forward (downward) motion is truncated to the distance to the nearest trap. If the no-space mode is on, no spacing occurs (see ns , and rs below).
.sv <i>N</i>	-	<i>N=1 V</i>	v	Save a contiguous vertical block of size <i>N</i> . If the distance to the next trap is greater than <i>N</i> , <i>N</i> vertical space is output. No-space mode has <i>no</i> effect. If this distance is less than <i>N</i> , no vertical space is immediately output, but <i>N</i> is remembered for later output (see os). Subsequent sv requests will overwrite any still remembered <i>N</i> .
.os	-	-	-	Output saved vertical space. No-space mode has <i>no</i> effect. Used to finally output a block of vertical space requested by an earlier sv request.
.ns	space	-	D	No-space mode turned on. When on, the no-space mode inhibits sp requests and bp requests <i>without</i> a next page number. The no-space mode is turned off when a line of output occurs, or with rs .
.rs	space	-	D	Restore spacing. The no-space mode is turned off.
Blank text line.	-	-	B	Causes a break and output of a blank line exactly like sp 1 .

6. Line Length and Indenting

The maximum line length for fill mode may be set with **ll**. The indent may be set with **in**; an indent applicable to *only* the *next* output line may be set with **ti**. The line length includes indent space but *not* page offset space. The line-length minus the indent is the basis for centering with **ce**. The effect of **ll**, **in**, or **ti** is delayed, if a partially collected line exists, until after that line is output. In fill mode the length of text on an output line is less than or equal to the line length minus the indent. The current line length and indent are available in registers **.l** and **.i** respectively. The length of *three-part titles* produced by **tl** (see §14) is *independently* set by **lt**.

Request Form	Initial Value	If No Argument	Notes	Explanation
.ll $\pm N$	6.5 in	previous	E,m	Line length is set to $\pm N$. In TROFF the maximum (line-length) + (page-offset) is about 7.54 inches.
.in $\pm N$	<i>N=0</i>	previous	B,E,m	Indent is set to $\pm N$. The indent is prepended to each output line.
.ti $\pm N$	-	ignored	B,E,m	Temporary indent. The <i>next</i> output text line will be indented a distance $\pm N$ with respect to the current indent. The resulting total indent may not be negative. The current indent is not changed.

7. Macros, Strings, Diversion, and Position Traps

7.1. Macros and strings. A *macro* is a named set of arbitrary *lines* that may be invoked by name or with a *trap*. A *string* is a named string of *characters*, *not* including a newline character, that may be interpolated by name at any point. Request, macro, and string names share the *same* name list. Macro and string names may be one or two characters long and may usurp previously defined request, macro, or string names. Any of these entities may be renamed with **rn** or removed with **rm**. Macros are created by **de** and **di**, and appended to by **am** and **da**; **di** and **da** cause normal output to be stored in a macro. Strings are created by **ds** and appended to by **as**. A macro is invoked in the same way as a request; a

control line beginning `.xx` will interpolate the contents of macro `xx`. The remainder of the line may contain up to nine *arguments*. The strings `x` and `xx` are interpolated at any desired point with `*x` and `*(xx` respectively. String references and macro invocations may be nested.

7.2. *Copy mode input interpretation.* During the definition and extension of strings and macros (not by diversion) the input is read in *copy mode*. The input is copied without interpretation *except* that:

- The contents of number registers indicated by `\n` are interpolated.
- Strings indicated by `*` are interpolated.
- Arguments indicated by `\$` are interpolated.
- Concealed newlines indicated by `\(newline)` are eliminated.
- Comments indicated by `\#` are eliminated.
- `\t` and `\a` are interpreted as ASCII horizontal tab and SOH respectively (§9).
- `\\` is interpreted as `\`.
- `\.` is interpreted as `."`.

These interpretations can be suppressed by prepending a `\`. For example, since `\\` maps into a `\`, `\\n` will copy as `\n` which will be interpreted as a number register indicator when the macro or string is reread.

7.3. *Arguments.* When a macro is invoked by name, the remainder of the line is taken to contain up to nine arguments. The argument separator is the space character, and arguments may be surrounded by double-quotes to permit imbedded space characters. Pairs of double-quotes may be imbedded in double-quoted arguments to represent a single double-quote. If the desired arguments won't fit on a line, a concealed newline may be used to continue on the next line.

When a macro is invoked the *input level* is *pushed down* and any arguments available at the previous level become unavailable until the macro is completely read and the previous level is restored. A macro's own arguments can be interpolated at *any* point within the macro with `\$N`, which interpolates the *N*th argument ($1 \leq N \leq 9$). If an invoked argument doesn't exist, a null string results. For example, the macro `xx` may be defined by

```
.de xx      \"begin definition
Today is \\$1 the \\$2.
..         \"end definition
```

and called by

```
.xx Monday 14th
```

to produce the text

```
Today is Monday the 14th.
```

Note that the `\$` was concealed in the definition with a prepended `\`. The number of currently available arguments is in the `.$` register.

No arguments are available at the top (non-macro) level in this implementation. Because string referencing is implemented as a input-level push down, no arguments are available from *within* a string. No arguments are available within a trap-invoked macro.

Arguments are copied in *copy mode* onto a stack where they are available for reference. The mechanism does not allow an argument to contain a direct reference to a *long* string (interpolated at copy time) and it is advisable to conceal string references (with an extra `\`) to delay interpolation until argument reference time.

7.4. *Diversions.* Processed output may be diverted into a macro for purposes such as footnote processing (see Tutorial §T5) or determining the horizontal and vertical size of some text for conditional changing of pages or columns. A single diversion trap may be set at a specified vertical position. The number registers `dn` and `dl` respectively contain the vertical and horizontal size of the most recently ended diversion. Processed text that is diverted into a macro retains the vertical size of each of its lines when reread in *nofill* mode regardless of the current *V*. Constant-spaced (`cs`) or emboldened (`bd`) text that is diverted can be reread correctly only if these modes are again or still in effect at reread time. One way

to do this is to imbed in the diversion the appropriate **cs** or **bd** requests with the *transparent* mechanism described in §10.6.

Diversions may be nested and certain parameters and registers are associated with the current diversion level (the top non-diversion level may be thought of as the 0th diversion level). These are the diversion trap and associated macro, no-space mode, the internally-saved marked place (see **mk** and **rt**), the current vertical place (**.d** register), the current high-water text base-line (**.h** register), and the current diversion name (**.z** register).

7.5. Traps. Three types of trap mechanisms are available—page traps, a diversion trap, and an input-line-count trap. Macro-invocation traps may be planted using **wh** at any page position including the top. This trap position may be changed using **ch**. Trap positions at or below the bottom of the page have no effect unless or until moved to within the page or rendered effective by an increase in page length. Two traps may be planted at the *same* position only by first planting them at different positions and then moving one of the traps; the first planted trap will conceal the second unless and until the first one is moved (see Tutorial Examples §T5). If the first one is moved back, it again conceals the second trap. The macro associated with a page trap is automatically invoked when a line of text is output whose vertical size *reaches* or *sweeps past* the trap position. Reaching the bottom of a page springs the top-of-page trap, if any, provided there is a next page. The distance to the next trap position is available in the **.t** register; if there are no traps between the current position and the bottom of the page, the distance returned is the distance to the page bottom.

A macro-invocation trap effective in the current diversion may be planted using **dt**. The **.t** register works in a diversion; if there is no subsequent trap a *large* distance is returned. For a description of input-line-count traps, see it below.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.de <i>xx yy</i>	-	<i>.yy=..</i>	-	Define or redefine the macro <i>xx</i> . The contents of the macro begin on the next input line. Input lines are copied in <i>copy mode</i> until the definition is terminated by a line beginning with <i>.yy</i> , whereupon the macro <i>yy</i> is called. In the absence of <i>yy</i> , the definition is terminated by a line beginning with <i>..</i> . A macro may contain de requests provided the terminating macros differ or the contained definition terminator is concealed. <i>..</i> can be concealed as <i>\\.</i> which will copy as <i>\\.</i> and be reread as <i>..</i> .
.am <i>xx yy</i>	-	<i>.yy=..</i>	-	Append to macro (append version of de).
.ds <i>xx string</i>	-	ignored	-	Define a string <i>xx</i> containing <i>string</i> . Any initial double-quote in <i>string</i> is stripped off to permit initial blanks.
.as <i>xx string</i>	-	ignored	-	Append <i>string</i> to string <i>xx</i> (append version of ds).
.rm <i>xx</i>	-	ignored	-	Remove request, macro, or string. The name <i>xx</i> is removed from the name list and any related storage space is freed. Subsequent references will have no effect.
.rn <i>xx yy</i>	-	ignored	-	Rename request, macro, or string <i>xx</i> to <i>yy</i> . If <i>yy</i> exists, it is first removed.
.di <i>xx</i>	-	end	D	Divert output to macro <i>xx</i> . Normal text processing occurs during diversion except that page offsetting is not done. The diversion ends when the request di or da is encountered without an argument; extraneous requests of this type should not appear when nested diversions are being used.

.da <i>xx</i>	-	end	D	Divert, appending to <i>xx</i> (append version of df).
.wh <i>N xx</i>	-	-	v	Install a trap to invoke <i>xx</i> at page position <i>N</i> ; a <i>negative N</i> will be interpreted with respect to the page bottom. Any macro previously planted at <i>N</i> is replaced by <i>xx</i> . A zero <i>N</i> refers to the top of a page. In the absence of <i>xx</i> , the first found trap at <i>N</i> , if any, is removed.
.ch <i>xx N</i>	-	-	v	Change the trap position for macro <i>xx</i> to be <i>N</i> . In the absence of <i>N</i> , the trap, if any, is removed.
.dt <i>N xx</i>	-	off	D,v	Install a diversion trap at position <i>N</i> in the <i>current</i> diversion to invoke macro <i>xx</i> . Another dt will redefine the diversion trap. If no arguments are given, the diversion trap is removed.
.it <i>N xx</i>	-	off	E	Set an input-line-count trap to invoke the macro <i>xx</i> after <i>N</i> lines of <i>text</i> input have been read (control or request lines don't count). The text may be in-line text or text interpolated by inline or trap-invoked macros.
.em <i>xx</i>	none	none	-	The macro <i>xx</i> will be invoked when all input has ended. The effect is the same as if the contents of <i>xx</i> had been at the end of the last file processed.

8. Number Registers

A variety of parameters are available to the user as predefined, named *number registers* (see Summary and Index, page 7). In addition, the user may define his own named registers. Register names are one or two characters long and *do not* conflict with request, macro, or string names. Except for certain predefined read-only registers, a number register can be read, written, automatically incremented or decremented, and interpolated into the input in a variety of formats. One common use of user-defined registers is to automatically number sections, paragraphs, lines, etc. A number register may be used any time numerical input is expected or desired and may be used in numerical *expressions* (§1.4).

Number registers are created and modified using **nr**, which specifies the name, numerical value, and the auto-increment size. Registers are also modified, if accessed with an auto-incrementing sequence. If the registers *x* and *xx* both contain *N* and have the auto-increment size *M*, the following access sequences have the effect shown:

Sequence	Effect on Register	Value Interpolated
$\backslash n x$	none	<i>N</i>
$\backslash n (xx)$	none	<i>N</i>
$\backslash n + x$	<i>x</i> incremented by <i>M</i>	<i>N + M</i>
$\backslash n - x$	<i>x</i> decremented by <i>M</i>	<i>N - M</i>
$\backslash n + (xx)$	<i>xx</i> incremented by <i>M</i>	<i>N + M</i>
$\backslash n - (xx)$	<i>xx</i> decremented by <i>M</i>	<i>N - M</i>

When interpolated, a number register is converted to decimal (default), decimal with leading zeros, lower-case Roman, upper-case Roman, lower-case sequential alphabetic, or upper-case sequential alphabetic according to the format specified by **af**.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.nr <i>R ± N M</i>	-	-	u	The number register <i>R</i> is assigned the value $\pm N$ with respect to the previous value, if any. The increment for auto-incrementing is set to <i>M</i> .

.ar R c arabic - Assign format *c* to register *R*. The available formats are:

Format	Numbering Sequence
1	0,1,2,3,4,5,...
001	000,001,002,003,004,005,...
i	0,i,ii,iii,iv,v,...
I	0,I,II,III,IV,V,...
a	0,a,b,c, ..., z,aa,ab, ..., zz,aaa,...
A	0,A,B,C, ..., Z,AA,AB, ..., ZZ,AAA,...

An arabic format having *N* digits specifies a field width of *N* digits (example 2 above). The read-only registers and the *width* function (§11.2) are always arabic.

.rr R - ignored - Remove register *R*. If many registers are being created dynamically, it may become necessary to remove no longer used registers to recapture internal storage space for newer registers.

9. Tabs, Leaders, and Fields

9.1. Tabs and leaders. The ASCII horizontal tab character and the ASCII SOH (hereafter known as the *leader* character) can both be used to generate either horizontal motion or a string of repeated characters. The length of the generated entity is governed by internal *tab stops* specifiable with **ta**. The default difference is that tabs generate motion and leaders generate a string of periods; **tc** and **lc** offer the choice of repeated character or motion. There are three types of internal tab stops—*left* adjusting, *right* adjusting, and *centering*. In the following table: *D* is the distance from the current position on the *input* line (where a tab or leader was found) to the next tab stop; *next-string* consists of the input characters following the tab (or leader) up to the next tab (or leader) or end of line; and *W* is the width of *next-string*.

Tab type	Length of motion or repeated characters	Location of <i>next-string</i>
Left	<i>D</i>	Following <i>D</i>
Right	<i>D - W</i>	Right adjusted within <i>D</i>
Centered	<i>D - W/2</i>	Centered on right end of <i>D</i>

The length of generated motion is allowed to be negative, but that of a repeated character string cannot be. Repeated character strings contain an integer number of characters, and any residual distance is prepended as motion. Tabs or leaders found after the last tab stop are ignored, but may be used as *next-string* terminators.

Tabs and leaders are not interpreted in *copy mode*. **\t** and **\a** always generate a non-interpreted tab and leader respectively, and are equivalent to actual tabs and leaders in *copy mode*.

9.2. Fields. A *field* is contained between a pair of *field delimiter* characters, and consists of sub-strings separated by *padding* indicator characters. The field length is the distance on the *input* line from the position where the field begins to the next tab stop. The difference between the total length of all the sub-strings and the field length is incorporated as horizontal padding space that is divided among the indicated padding places. The incorporated padding is allowed to be negative. For example, if the field delimiter is **#** and the padding indicator is **^**, **#^xxx^right#** specifies a right-adjusted string with the string *xxx* centered in the remaining space.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.ta <i>N</i> ...	0.8; 0.5in	none	E,m	Set tab stops and types. <i>t</i> =R, right adjusting; <i>t</i> =C, centering; <i>t</i> absent, left adjusting. TROFF tab stops are preset every 0.5in.; NROFF every 0.8in. The stop values are separated by spaces, and a value preceded by + is treated as an increment to the previous stop value.
.tc <i>c</i>	none	none	E	The tab repetition character becomes <i>c</i> , or is removed specifying motion.
.lc <i>c</i>	.	none	E	The leader repetition character becomes <i>c</i> , or is removed specifying motion.
.fc <i>a b</i>	off	off	-	The field delimiter is set to <i>a</i> ; the padding indicator is set to the <i>space</i> character or to <i>b</i> , if given. In the absence of arguments the field mechanism is turned off.

10. Input and Output Conventions and Character Translations

10.1. Input character translations. Ways of inputting the graphic character set were discussed in §2.1. The ASCII control characters horizontal tab (§9.1), SOH (§9.1), and backspace (§10.3) are discussed elsewhere. The newline delimits input lines. In addition, STX, ETX, ENQ, ACK, and BEL are accepted, and may be used as delimiters or translated into a graphic with **tr** (§10.5). *All others are ignored.*

The *escape* character **** introduces *escape sequences*—causes the following character to mean another character, or to indicate some function. A complete list of such sequences is given in the Summary and Index on page 6. **** should not be confused with the ASCII control character ESC of the same name. The escape character **** can be input with the sequence ****. The escape character can be changed with **ec**, and all that has been said about the default **** becomes true for the new escape character. **\e** can be used to print whatever the current escape character is. If necessary or convenient, the escape mechanism may be turned off with **eo**, and restored with **ec**.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.ec <i>c</i>	\	\	-	Set escape character to \ , or to <i>c</i> , if given.
.eo	on	-	-	Turn escape mechanism off.

10.2. Ligatures. Five ligatures are available in the current TROFF character set — **fi**, **fl**, **ff**, **ffi**, and **ffl**. They may be input (even in NROFF) by **\(fi**, **\(fl**, **\(ff**, **\(ffi**, and **\(ffl** respectively. The ligature mode is normally on in TROFF, and *automatically* invokes ligatures during input.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.lg <i>N</i>	off; on	on	-	Ligature mode is turned on if <i>N</i> is absent or non-zero, and turned off if <i>N</i> =0. If <i>N</i> =2, only the two-character ligatures are automatically invoked. Ligature mode is inhibited for request, macro, string, register, or file names, and in <i>copy mode</i> . No effect in NROFF.

10.3. Backspacing, underlining, overstriking, etc. Unless in *copy mode*, the ASCII backspace character is replaced by a backward horizontal motion having the width of the space character. Underlining as a form of line-drawing is discussed in §12.4. A generalized overstriking function is described in §12.1.

NROFF automatically underlines characters in the *underline* font, specifiable with **uf**, normally that on font position 2 (normally Times Italic, see §2.2). In addition to **ft** and **\fF**, the underline font may be selected by **ul** and **cu**. Underlining is restricted to an output-device-dependent subset of *reasonable* characters.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.ul <i>N</i>	off	<i>N=1</i>	E	Underline in NROFF (italicize in TROFF) the next <i>N</i> input text lines. Actually, switch to <i>underline</i> font, saving the current font for later restoration; <i>other</i> font changes within the span of a ul will take effect, but the restoration will undo the last change. Output generated by tl (§14) is affected by the font change, but does <i>not</i> decrement <i>N</i> . If <i>N</i> > 1, there is the risk that a trap interpolated macro may provide text lines within the span; environment switching can prevent this.
.cu <i>N</i>	off	<i>N=1</i>	E	A variant of ul that causes <i>every</i> character to be underlined in NROFF. Identical to ul in TROFF.
.uf <i>F</i>	Italic	Italic	-	Underline font set to <i>F</i> . In NROFF, <i>F</i> may <i>not</i> be on position 1 (initially Times Roman).

10.4. *Control characters.* Both the control character **.** and the *no-break* control character **'** may be changed, if desired. Such a change must be compatible with the design of any macros used in the span of the change, and particularly of any trap-invoked macros.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.cc <i>c</i>	.	.	E	The basic control character is set to <i>c</i> , or reset to ".".
.c2 <i>c</i>	'	'	E	The <i>nobreak</i> control character is set to <i>c</i> , or reset to "'".

10.5. *Output translation.* One character can be made a stand-in for another character using **tr**. All text processing (e. g. character comparisons) takes place with the input (stand-in) character which appears to have the width of the final character. The graphic translation occurs at the moment of output (including diversion).

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.tr <i>abcd....</i>	none	-	O	Translate <i>a</i> into <i>b</i> , <i>c</i> into <i>d</i> , etc. If an odd number of characters is given, the last one will be mapped into the space character. To be consistent, a particular translation must stay in effect from <i>input</i> to <i>output</i> time.

10.6. *Transparent throughput.* An input line beginning with a **\!** is read in *copy mode* and *transparently* output (without the initial **\!**); the text processor is otherwise unaware of the line's presence. This mechanism may be used to pass control information to a post-processor or to imbed control lines in a macro created by a diversion.

10.7. *Comments and concealed newlines.* An uncomfortably long input line that must stay one line (e. g. a string definition, or nofilled text) can be split into many physical lines by ending all but the last one with the escape ****. The sequence **\(newline)** is *always* ignored—except in a comment. Comments may be imbedded at the *end* of any line by prefacing them with ****". The newline at the end of a comment cannot be concealed. A line beginning with ****" will appear as a blank line and behave like **.sp 1**; a comment can be on a line by itself by beginning the line with ****".

11. Local Horizontal and Vertical Motions, and the Width Function

11.1. *Local Motions.* The functions **\v'N'** and **\h'N'** can be used for *local* vertical and horizontal motion respectively. The distance *N* may be negative; the *positive* directions are *rightward* and *downward*. A *local* motion is one contained *within* a line. To avoid unexpected vertical dislocations, it is necessary that the *net* vertical local motion within a word in filled text and otherwise within a line balance to zero. The above and certain other escape sequences providing local motion are summarized in the following table.

Vertical Local Motion	Effect in		Horizontal Local Motion	Effect in	
	TROFF	NROFF		TROFF	NROFF
$\backslash v' N'$	Move distance N		$\backslash h' N'$	Move distance N	
			$\backslash (\text{space})$	Unpaddable space-size space	
			$\backslash 0$	Digit-size space	
$\backslash u$	1/2 em up	1/2 line up	$\backslash $ $\backslash ^$	1/6 em space	ignored
$\backslash d$	1/2 em down	1/2 line down		1/12 em space	ignored
$\backslash r$	1 em up	1 line up			

As an example, E^2 could be generated by the sequence $E\backslash s-2\backslash v'-0.4m'2\backslash v'0.4m'\backslash s+2$; it should be noted in this example that the 0.4 em vertical motions are at the smaller size.

11.2. *Width Function.* The *width* function $\backslash w' string'$ generates the numerical width of *string* (in basic units). Size and font changes may be safely imbedded in *string*, and will not affect the current environment. For example, $\backslash ti-\backslash w'1. 'u$ could be used to temporarily indent leftward a distance equal to the size of the string "1. ".

The width function also sets three number registers. The registers **st** and **sb** are set respectively to the highest and lowest extent of *string* relative to the baseline; then, for example, the total *height* of the string is $\backslash n(stu-\backslash n(sbu$. In TROFF the number register **ct** is set to a value between 0 and 3: 0 means that all of the characters in *string* were short lower case characters without descenders (like **e**); 1 means that at least one character has a descender (like **y**); 2 means that at least one character is tall (like **H**); and 3 means that both tall characters and characters with descenders are present.

11.3. *Mark horizontal place.* The escape sequence $\backslash kx$ will cause the *current* horizontal position in the *input line* to be stored in register *x*. As an example, the construction $\backslash kx word\backslash h'\backslash nxu+2u' word$ will embolden *word* by backing up to almost its beginning and overprinting it, resulting in **word**.

12. Overstrike, Bracket, Line-drawing, and Zero-width Functions

12.1. *Overstriking.* Automatically centered overstriking of up to nine characters is provided by the *overstrike* function $\backslash o' string'$. The characters in *string* overprinted with centers aligned; the total width is that of the widest character. *string* should *not* contain local vertical motion. As examples, $\backslash o'e''$ produces \acute{e} , and $\backslash o'(mo)\backslash sl'$ produces £ .

12.2. *Zero-width characters.* The function $\backslash zc$ will output *c* without spacing over it, and can be used to produce left-aligned overstruck combinations. As examples, $\backslash z(ci)\backslash pl$ will produce ⊕ , and $\backslash br\backslash z(rn)\backslash ul\backslash br$ will produce the smallest possible constructed box \square .

12.3. *Large Brackets.* The Special Mathematical Font contains a number of bracket construction pieces ($\{ \} [] \{ \} [\} [\} [\}$) that can be combined into various bracket styles. The function $\backslash b' string'$ may be used to pile up vertically the characters in *string* (the first character on top and the last at the bottom); the characters are vertically separated by 1 em and the total pile is centered 1/2 em above the current baseline (1/2 line in NROFF). For example, $\backslash b'\backslash lc\backslash lf'E'\backslash b'\backslash rc\backslash rf'\backslash x'-0.5m'\backslash x'0.5m'$ produces $\left[E \right]$.

12.4. *Line drawing.* The function $\backslash l' Nc'$ will draw a string of repeated *c*'s towards the right for a distance N . ($\backslash l$ is \backslash (lower case L). If *c* looks like a continuation of an expression for N , it may insulated from N with a $\backslash \&$. If *c* is not specified, the $_$ (baseline rule) is used (underline character in NROFF). If N is negative, a backward horizontal motion of size N is made *before* drawing the string. Any space resulting from $N/(\text{size of } c)$ having a remainder is put at the beginning (left end) of the string. In the case of characters that are designed to be connected such as baseline-rule $_$, underrule $_$, and root-en $_$, the remainder space is covered by over-lapping. If N is *less* than the width of *c*, a single *c* is centered on a distance N . As an example, a macro to underscore a string can be written

```
.de us
\ \ $! \ l' | 0 \ ul'
..
```

or one to draw a box around a string

```
.de bx
\{br\|\\$1\|\\(br\|'|0\{rn\|'|0\{ul'
..
```

such that

```
.ul "underlined words"
```

and

```
.bx "words in a box"
```

yield underlined words and words in a box.

The function `\L' Nc'` will draw a vertical line consisting of the (optional) character *c* stacked vertically apart 1 em (1 line in NROFF), with the first two characters overlapped, if necessary, to form a continuous line. The default character is the *box rule* | (`\(br`); the other suitable character is the *bold vertical* | (`\(bv`). The line is begun without any initial motion relative to the current base line. A positive *N* specifies a line drawn downward and a negative *N* specifies a line drawn upward. After the line is drawn no compensating motions are made; the instantaneous baseline is at the *end* of the line.

The horizontal and vertical line drawing functions may be used in combination to produce large boxes. The zero-width *box-rule* and the 1/2-em wide *underrule* were *designed* to form corners when using 1-em vertical spacings. For example the macro

```
.de eb
.sp -1      \ "compensate for next automatic base-line spacing
.nf        \ "avoid possibly overflowing word buffer
\h'-.5n\L'\|\\nau-1'\|\\n(.lu+1n\{ul'\L'-|\\nau+1'\|0u-.5n\{ul'  \ "draw box
.fi
..
```

will draw a box around some text whose beginning vertical place was saved in number register *a* (e. g. using `.mk a`) as done for this paragraph.

13. Hyphenation.

The automatic hyphenation may be switched off and on. When switched on with `hy`, several variants may be set. A *hyphenation indicator* character may be imbedded in a word to specify desired hyphenation points, or may be prepended to suppress hyphenation. In addition, the user may specify a small exception word list.

Only words that consist of a central alphabetic string surrounded by (usually null) non-alphabetic strings are considered candidates for automatic hyphenation. Words that were input containing hyphens (minus), em-dashes (`\(em`), or hyphenation indicator characters—such as mother-in-law—are *always* subject to splitting after those characters, whether or not automatic hyphenation is on or off.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
<code>.nh</code>	hyphenate	-	F	Automatic hyphenation is turned off.
<code>.hyN</code>	on, N=1	on, N=1	E	Automatic hyphenation is turned on for $N \geq 1$, or off for $N=0$. If $N=2$, <i>last</i> lines (ones that will cause a trap) are not hyphenated. For $N=4$ and 8, the last and first two characters respectively of a word are not split off. These values are additive; i. e. $N=14$ will invoke all three restrictions.
<code>.hc c</code>	\%	\%	E	Hyphenation: indicator character is set to <i>c</i> or to the default \%. The indicator does not appear in the output.
<code>.hw wordl ...</code>		ignored	-	Specify hyphenation points in words with imbedded minus signs. Versions of a word with terminal <i>s</i> are

implied; i. e. *dig - n* implies *dig - ns*. This list is examined initially *and* after each suffix stripping. The space available is small—about 128 characters.

14. Three Part Titles.

The titling function **tl** provides for automatic placement of three fields at the left, center, and right of a line with a title-length specifiable with **lt**. **tl** may be used anywhere, and is independent of the normal text collecting process. A common use is in header and footer macros.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
<code>.tl 'left' center' right'</code>	-	-	-	The strings <i>left</i> , <i>center</i> , and <i>right</i> are respectively left-adjusted, centered, and right-adjusted in the current title-length. Any of the strings may be empty, and overlapping is permitted. If the page-number character (initially %) is found within any of the fields it is replaced by the current page number having the format assigned to register %. Any character may be used as the string delimiter.
<code>.pc c</code>	%	off		The page number character is set to <i>c</i> , or removed. The page-number register remains %.
<code>.lt ±N</code>	6.5 in	previous	E,m	Length of title set to ± <i>N</i> . The line-length and the title-length are <i>independent</i> . Indents do not apply to titles; page-offsets do.

15. Output Line Numbering.

Automatic sequence numbering of output lines may be requested with **nm**. When in effect, a three-digit, arabic number plus a digit-space is prepended to output text lines. The text lines are 3 thus offset by four digit-spaces, and otherwise retain their line length; a reduction in line length may be desired to keep the right margin aligned with an earlier margin. Blank lines, other vertical spaces, and lines generated by **tl** are *not* numbered. Numbering can be temporarily suspended with 6 **nn**, or with an **.nm** followed by a later **.nm +0**. In addition, a line number indent *I*, and the number-text separation *S* may be specified in digit-spaces. Further, it can be specified that only those line numbers that are multiples of some number *M* are to be printed (the others will appear 9 as blank number fields).

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
<code>.nm ±N M S I</code>		off	E	Line number mode. If ± <i>N</i> is given, line numbering is turned on, and the next output line numbered is numbered ± <i>N</i> . Default values are <i>M</i> =1, <i>S</i> =1, and <i>I</i> =0. Parameters corresponding to missing arguments are unaffected; a non-numeric argument is considered missing. In the absence of all arguments, numbering is turned off; the next line number is preserved for possible further use in number register ln .
<code>.nn N</code>	-	<i>N</i> =1	E	The next <i>N</i> text output lines are not numbered.

As an example, the paragraph portions of this section are numbered with *M*=3: **.nm 1 3** was placed at the beginning; **.nm** was placed at the end of the first paragraph; and **.nm +0** was placed 12 in front of this paragraph; and **.nm** finally placed at the end. Line lengths were also changed (by `\w'0000'u`) to keep the right side aligned. Another example is **.nm +5 5 x 3** which turns on numbering with the line number of the next line to be 5 greater than the last numbered line, with 15 *M*=5, with spacing *S* untouched, and with the indent *I* set to 3.

16. Conditional Acceptance of Input

In the following, *c* is a one-character, built-in *condition* name, *!* signifies *not*, *N* is a numerical expression, *string1* and *string2* are strings delimited by any non-blank, non-numeric character *not* in the strings, and *anything* represents what is conditionally accepted.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
<code>.if c anything</code>	-	-	-	If condition <i>c</i> true, accept <i>anything</i> as input; in multi-line case use <code>\{anything\}</code> .
<code>.if !c anything</code>	-	-	-	If condition <i>c</i> false, accept <i>anything</i> .
<code>.if N anything</code>	-	-	u	If expression $N > 0$, accept <i>anything</i> .
<code>.if !N anything</code>	-	-	u	If expression $N \leq 0$, accept <i>anything</i> .
<code>.if 'string1' string2' anything</code>	-	-	-	If <i>string1</i> identical to <i>string2</i> , accept <i>anything</i> .
<code>.if !'string1' string2' anything</code>	-	-	-	If <i>string1</i> not identical to <i>string2</i> , accept <i>anything</i> .
<code>.ie c anything</code>	-	-	u	If portion of if-else; all above forms (like <code>if</code>).
<code>.el anything</code>	-	-	-	Else portion of if-else.

The built-in condition names are:

Condition Name	True If
o	Current page number is odd
e	Current page number is even
t	Formatter is TROFF
n	Formatter is NROFF

If the condition *c* is *true*, or if the number *N* is greater than zero, or if the strings compare identically (including motions and character size and font), *anything* is accepted as input. If a *!* precedes the condition, number, or string comparison, the sense of the acceptance is reversed.

Any spaces between the condition and the beginning of *anything* are skipped over. The *anything* can be either a single input line (text, macro, or whatever) or a number of input lines. In the multi-line case, the first line must begin with a left delimiter `\{` and the last line must end with a right delimiter `\}`.

The request `ie` (if-else) is identical to `if` except that the acceptance state is remembered. A subsequent and matching `el` (else) request then uses the reverse sense of that state. `ie - el` pairs may be nested.

Some examples are:

```
.if e .tl 'Even Page %''
```

which outputs a title if the page number is even; and

```
.ie \n% > 1 \{\
.sp 0.5i
.tl 'Page %''
.sp |1.2i \}
.el .sp |2.5i
```

which treats page 1 differently from other pages.

17. Environment Switching.

A number of the parameters that control the text processing are gathered together into an *environment*, which can be switched by the user. The environment parameters are those associated with requests noting **E** in their *Notes* column; in addition, partially collected lines and words are in the environment. Everything else is global; examples are page-oriented parameters, diversion-oriented parameters,

number registers, and macro and string definitions. All environments are initialized with default parameter values.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.ev <i>N</i>	<i>N=0</i>	previous	-	Environment switched to environment $0 \leq N \leq 2$. Switching is done in push-down fashion so that restoring a previous environment <i>must</i> be done with .ev rather than specific reference.

18. Insertions from the Standard Input

The input can be temporarily switched to the system *standard input* with **rd**, which will switch back when *two* newlines in a row are found (the *extra* blank line is not used). This mechanism is intended for insertions in form-letter-like documentation. On UNIX, the *standard input* can be the user's keyboard, a *pipe*, or a *file*.

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.rd <i>prompt</i>	-	<i>prompt=BEL</i>	-	Read insertion from the standard input until two newlines in a row are found. If the standard input is the user's keyboard, <i>prompt</i> (or a BEL) is written onto the user's terminal. rd behaves like a macro, and arguments may be placed after <i>prompt</i> .
.ex	-	-	-	Exit from NROFF/TROFF. Text processing is terminated exactly as if all input had ended.

If insertions are to be taken from the terminal keyboard *while* output is being printed on the terminal, the command line option **-q** will turn off the echoing of keyboard input and prompt only with BEL. The regular input and insertion input *cannot* simultaneously come from the standard input.

As an example, multiple copies of a form letter may be prepared by entering the insertions for all the copies in one file to be used as the standard input, and causing the file containing the letter to reinvoke itself using **nx** (§19); the process would ultimately be ended by an **ex** in the insertion file.

19. Input/Output File Switching

<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.so <i>filename</i>	-	-	-	Switch source file. The top input (file reading) level is switched to <i>filename</i> . The effect of an so encountered in a macro is not felt until the input level returns to the file level. When the new file ends, input is again taken from the original file. so 's may be nested.
.nx <i>filename</i>	-	end-of-file	-	Next file is <i>filename</i> . The current file is considered ended, and the input is immediately switched to <i>filename</i> .
.pi <i>program</i>	-	-	-	Pipe output to <i>program</i> (NROFF only). This request must occur <i>before</i> any printing occurs. No arguments are transmitted to <i>program</i> .

20. Miscellaneous

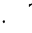
<i>Request Form</i>	<i>Initial Value</i>	<i>If No Argument</i>	<i>Notes</i>	<i>Explanation</i>
.mc <i>c N</i>	-	off	E,m	Specifies that a <i>margin</i> character <i>c</i> appear a distance <i>N</i> to the right of the right margin after each non-empty text line (except those produced by tl). If the output line is too-long (as can happen in nofill mode) the character will

be appended to the line. If *N* is not given, the previous *N* is used; the initial *N* is 0.2 inches in NROFF and 1 cm in TROFF. The margin character used with this paragraph was a 12-point box-rule.

.tm	<i>string</i>	-	newline	-	After skipping initial blanks, <i>string</i> (rest of the line) is read in <i>copy mode</i> and written on the user's terminal.
.ig	<i>yy</i>	-	<i>.yy=..</i>	-	Ignore input lines. ig behaves exactly like de (§7) except that the input is discarded. The input is read in <i>copy mode</i> , and any auto-incremented registers will be affected.
.pm	<i>t</i>	-	all	-	Print macros. The names and sizes of all of the defined macros and strings are printed on the user's terminal; if <i>t</i> is given, only the total of the sizes is printed. The sizes is given in <i>blocks</i> of 128 characters.
.fl		-		B	Flush output buffer. Used in interactive debugging to force output.

21. Output and Error Messages.

The output from **tm**, **pm**, and the prompt from **rd**, as well as various *error* messages are written onto UNIX's *standard message* output. The latter is different from the *standard output*, where NROFF formatted output goes. By default, both are written onto the user's terminal, but they can be independently redirected.

Various *error* conditions may occur during the operation of NROFF and TROFF. Certain less serious errors having only local impact do not cause processing to terminate. Two examples are *word overflow*, caused by a word that is too large to fit into the word buffer (in fill mode), and *line overflow*, caused by an output line that grew too large to fit in the line buffer; in both cases, a message is printed, the offending excess is discarded, and the affected word or line is marked at the point of truncation with a * in NROFF and a  in TROFF. The philosophy is to continue processing, if possible, on the grounds that output useful for debugging may be produced. If a serious error occurs, processing terminates, and an appropriate message is printed. Examples are the inability to create, read, or write files, and the exceeding of certain internal limits that make future output unlikely to be useful.

TUTORIAL EXAMPLES

T1. Introduction

Although NROFF and TROFF have by design a syntax reminiscent of earlier text processors* with the intent of easing their use, it is almost always necessary to prepare at least a small set of macro definitions to describe most documents. Such common formatting needs as page margins and footnotes are deliberately not built into NROFF and TROFF. Instead, the macro and string definition, number register, diversion, environment switching, page-position trap, and conditional input mechanisms provide the basis for user-defined implementations.

The examples to be discussed are intended to be useful and somewhat realistic, but won't necessarily cover all relevant contingencies. Explicit numerical parameters are used in the examples to make them easier to read and to illustrate typical values. In many cases, number registers would really be used to reduce the number of places where numerical information is kept, and to concentrate conditional parameter initialization like that which depends on whether TROFF or NROFF is being used.

T2. Page Margins

As discussed in §3, *header* and *footer* macros are usually defined to describe the top and bottom page margin areas respectively. A trap is planted at page position 0 for the header, and at $-N$ (N from the page bottom) for the footer. The simplest such definitions might be

```
.de hd          \ "define header
`sp 1i
..             \ "end definition
.de fo          \ "define footer
`bp
..            \ "end definition
.wh 0 hd
.wh -1i fo
```

which provide blank 1 inch top and bottom margins. The header will occur on the *first* page, only if the definition and trap exist prior to the

*For example: P. A. Crisman, Ed., *The Compatible Time-Sharing System*, MIT Press, 1965, Section AH9.01 (Description of RUNOFF program on MIT's CTSS system).

initial pseudo-page transition (§3). In fill mode, the output line that springs the footer trap was typically forced out because some part or whole word didn't fit on it. If anything in the footer and header that follows causes a *break*, that word or part word will be forced out. In this and other examples, requests like **bp** and **sp** that normally cause breaks are invoked using the *no-break* control character ` to avoid this. When the header/footer design contains material requiring independent text processing, the environment may be switched, avoiding most interaction with the running text.

A more realistic example would be

```
.de hd          \ "header
.if t .tl `(rn`\ (rn` \ "troff cut mark
.if \n%>1 \{\
`sp|0.5i-1     \ "tl base at 0.5i
.tl "- % -"   \ "centered page number
.ps           \ "restore size
.ft          \ "restore font
.vs \}       \ "restore vs
`sp|1.0i     \ "space to 1.0i
.ns         \ "turn on no-space mode
..
.de fo          \ "footer
.ps 10       \ "set footer/header size
.ft R        \ "set font
.vs 12p      \ "set base-line spacing
.if \n%=1 \{\
`sp|\n(.pu-0.5i-1 \ "tl base 0.5i up
.tl "- % -" \} \ "first page number
`bp
..
.wh 0 hd
.wh -1i fo
```

which sets the size, font, and base-line spacing for the header/footer material, and ultimately restores them. The material in this case is a page number at the bottom of the first page and at the top of the remaining pages. If TROFF is used, a *cut mark* is drawn in the form of *root-en*'s at each margin. The **sp**'s refer to absolute positions to avoid dependence on the base-line spacing. Another reason for this in the footer is that the footer is invoked by printing a line whose vertical spacing swept past the trap position by possibly as

much as the base-line spacing. The *no-space mode* is turned on at the end of **hd** to render ineffective accidental occurrences of **sp** at the top of the running text.

The above method of restoring size, font, etc. presupposes that such requests (that set *previous value*) are *not* used in the running text. A better scheme is save and restore both the current and previous values as shown for size in the following:

```
.de fo
.nr s1 \\n(.s  \\"current size
.ps
.nr s2 \\n(.s  \\"previous size
. ---        \\"rest of footer
..
.de hd
. ---        \\"header stuff
.ps \\n(s2    \\"restore previous size
.ps \\n(s1    \\"restore current size
..
```

Page numbers may be printed in the bottom margin by a separate macro triggered during the footer's page ejection:

```
.de bn      \\"bottom number
.tl "" - % - "" \\"centered page number
..
.wh -0.5i-1v bn \\"tl base 0.5i up
```

T3. Paragraphs and Headings

The housekeeping associated with starting a new paragraph should be collected in a paragraph macro that, for example, does the desired preparagraph spacing, forces the correct font, size, base-line spacing, and indent, checks that enough space remains for *more than one* line, and requests a temporary indent.

```
.de pg      \\"paragraph
.br         \\"break
.ft R      \\"force font,
.ps 10     \\"size,
.vs 12p    \\"spacing,
.in 0      \\"and indent
.sp 0.4    \\"prespace
.ne 1+\\n(.Vu \\"want more than 1 line
.ti 0.2i   \\"temp indent
..
```

The first break in **pg** will force out any previous partial lines, and must occur before the **vs**. The forcing of font, etc. is partly a defense against prior error and partly to permit things like section heading macros to set parameters only once.

The prespacing parameter is suitable for TROFF; a larger space, at least as big as the output device vertical resolution, would be more suitable in NROFF. The choice of remaining space to test for in the **ne** is the smallest amount greater than one line (the **.V** is the available vertical resolution).

A macro to automatically number section headings might look like:

```
.de sc      \\"section
. ---      \\"force font, etc.
.sp 0.4    \\"prespace
.ne 2.4+\\n(.Vu \\"want 2.4+ lines
.fi
\\n+S.
..
.nr S 0 1   \\"init S
```

The usage is **.sc**, followed by the section heading text, followed by **.pg**. The **ne** test value includes one line of heading, 0.4 line in the following **pg**, and one line of the paragraph text. A word consisting of the next section number and a period is produced to begin the heading line. The format of the number may be set by **af** (§8).

Another common form is the labeled, indented paragraph, where the label protrudes left into the indent space.

```
.de lp      \\"labeled paragraph
.pg
.in 0.5i    \\"paragraph indent
.ta 0.2i 0.5i \\"label, paragraph
.ti 0
\\t\\$1\\t\\c \\"flow into paragraph
..
```

The intended usage is **.lp label**; *label* will begin at 0.2 inch, and cannot exceed a length of 0.3 inch without intruding into the paragraph. The label could be right adjusted against 0.4 inch by setting the tabs instead with **.ta 0.4iR 0.5i**. The last line of **lp** ends with **\c** so that it will become a part of the first line of the text that follows.

T4. Multiple Column Output

The production of multiple column pages requires the footer macro to decide whether it was invoked by other than the last column, so that it will begin a new column rather than produce the bottom margin. The header can initialize a column register that the footer will increment and test. The following is arranged for two columns, but is easily modified for more.

```
.de hd      \ "header
. ---
.nr cl 0 1  \ "init column count
.mk        \ "mark top of text
..
.de fo      \ "footer
.le \n+(cl<2 \{\
.po +3.4i  \ "next column; 3.1+0.3
.rt        \ "back to mark
.ns \}     \ "no-space mode
.el \{\
.po \nMu   \ "restore left margin
. ---
.bp \}
..
.ll 3.1i   \ "column width
.nr M \n(.o \ "save left margin
```

Typically a portion of the top of the first page contains full width text; the request for the narrower line length, as well as another `.mk` would be made where the two column output was to begin.

T5. Footnote Processing

The footnote mechanism to be described is used by imbedding the footnotes in the input text at the point of reference, demarcated by an initial `.fn` and a terminal `.ef`:

```
.fn
  Footnote text and control lines...
.ef
```

In the following, footnotes are processed in a separate environment and diverted for later printing in the space immediately prior to the bottom margin. There is provision for the case where the last collected footnote doesn't completely fit in the available space.

```
.de hd      \ "header
. ---
.nr x 0 1   \ "init footnote count
.nr y 0-\nb \ "current footer place
.ch fo -\nbu \ "reset footer trap
.if \n(dn .fz \ "leftover footnote
..
.de fo      \ "footer
.nr dn 0    \ "zero last diversion size
.if \nx \{\
.ev 1       \ "expand footnotes in ev 1
.nf        \ "retain vertical size
.FN        \ "footnotes
.rm FN     \ "delete it
.if "\n(.z"fy" .di \ "end overflow diversion
.nr x 0    \ "disable fx
```

```
.ev \}     \ "pop environment
. ---
.bp
..
.de fx      \ "process footnote overflow
.if \nx .di fy \ "divert overflow
..
.de fn      \ "start footnote
.da FN     \ "divert (append) footnote
.ev 1     \ "in environment 1
.if \n+x=1 .fs \ "if first, include separator
.fi        \ "fill mode
..
.de ef      \ "end footnote
.br        \ "finish output
.nr z \n(.v \ "save spacing
.ev        \ "pop ev
.di        \ "end diversion
.nr y -\n(dn \ "new footer position,
.if \nx=1 .nr y -(\n(.v-\nz) \
          \ "uncertainty correction
.ch fo \nyu \ "y is negative
.if (\n(nl+1v)>(\n(.p+\ny) \
.ch fo \n(nlu+1v \ "it didn't fit
..
.de fs      \ "separator
\l' 1i'    \ "1 inch rule
.br
..
.de fz      \ "get leftover footnote
.fn
.nf        \ "retain vertical size
.fy        \ "where fx put it
.ef
..
.nr b 1.0i  \ "bottom margin size
.wh 0 hd    \ "header trap
.wh 12i fo  \ "footer trap, temp position
.wh -\nbu fx \ "fx at footer position
.ch fo -\nbu \ "conceal fx with fo
```

The header `hd` initializes a footnote count register `x`, and sets both the current footer trap position register `y` and the footer trap itself to a nominal position specified in register `b`. In addition, if the register `dn` indicates a leftover footnote, `fz` is invoked to reprocess it. The footnote start macro `fn` begins a diversion (append) in environment 1, and increments the count `x`; if the count is one, the footnote separator `fs` is interpolated. The separator is kept in a separate macro to permit user redefinition. The footnote end macro `ef` restores the previous environment and ends the diversion after saving the spacing size in register `z`. `y` is then decremented by the size of the

footnote, available in **dn**; then on the first footnote, **y** is further decremented by the difference in **vertical** base-line spacings of the two environments, to prevent the late triggering the footer trap from causing the last line of the combined footnotes to overflow. The footer trap is then set to the lower (on the page) of **y** or the current page position (**nl**) plus one line, to allow for printing the reference line. If indicated by **x**, the footer **fo** rereads the footnotes from **FN** in nofill mode in environment 1, and deletes **FN**. If the footnotes were too large to fit, the macro **fx** will be trap-invoked to redirect the overflow into **fy**, and the register **dn** will later indicate to the header whether **fy** is empty. Both **fo** and **fx** are planted in the nominal footer trap position in an order that causes **fx** to be concealed unless the **fo** trap is moved. The footer then terminates the overflow diversion, if necessary, and zeros **x** to disable **fx**, because the uncertainty correction together with a not-too-late triggering of the footer can result in the footnote rereading finishing before reaching the **fx** trap.

A good exercise for the student is to combine the multiple-column and footnote mechanisms.

T6. The Last Page

After the last input file has ended, NROFF and TROFF invoke the *end macro* (§7), if any, and when it finishes, eject the remainder of the page. During the eject, any traps encountered are processed normally. At the *end* of this last page, processing terminates *unless* a partial line, word, or partial word remains. If it is desired that another page be started, the end-macro

```
.de en      \end-macro
\c
'bp
..
.em en
```

will deposit a null partial word, and effect another last page.

Table I

Font Style Examples

The following fonts are printed in 12-point, with a vertical spacing of 14-point, and with non-alphanumeric characters separated by 1/4 em space. The Special Mathematical Font was specially prepared for Bell Laboratories by Graphic Systems, Inc. of Hudson, New Hampshire. The Times Roman, Italic, and Bold are among the many standard font available from that company.

Times Roman

abcdefghijklmnopqrstuvwxy
ABCDEFGHIJKLMNPNPQRSTUVWXYZ
1234567890
! \$ % & () ' * + - . , / : ; = ? [] |
• □ - - _ 1/4 1/2 3/4 fi fl ff ffi ffl ° † ‡ © ®

Times Italic

abcdefghijklmnopqrstuvwxy
ABCDEFGHIJKLMNPNPQRSTUVWXYZ
1234567890
*! \$ % & () ' * + - . , / : ; = ? [] |*
• □ - - _ 1/4 1/2 3/4 fi fl ff ffi ffl ° † ‡ © ®

Times Bold

abcdefghijklmnopqrstuvwxy
ABCDEFGHIJKLMNPNPQRSTUVWXYZ
1234567890
! \$ % & () ' * + - . , / : ; = ? [] |
• ■ - - _ 1/4 1/2 3/4 fi fl ff ffi ffl ° † ‡ © ®

Special Mathematical Font

" ^ _ ` ~ / < > { } # @ + - = *
 $\alpha \beta \gamma \delta \epsilon \zeta \eta \theta \iota \kappa \lambda \mu \nu \xi \omicron \pi \rho \sigma \tau \upsilon \phi \chi \psi \omega$
 $\Gamma \Delta \Theta \Lambda \Xi \Pi \Sigma \Upsilon \Phi \Psi \Omega$
 $\sqrt{\quad} \geq \leq \equiv \sim \cong \neq \rightarrow \leftarrow \uparrow \downarrow \times \div \pm \cup \cap \subset \supset \subseteq \supseteq \infty \partial$
\$ \nabla \neg \int \alpha \emptyset \in \dagger \blacksquare \blacktriangleright \textcircled{a} | \circ () \{ \} [] [] []

Table II
Input Naming Conventions for ' , ` and -
and for Non-ASCII Special Characters

Non-ASCII characters and *minus* on the standard fonts.

<i>Input Character</i>	<i>Char Name</i>	<i>Character Name</i>	<i>Input Character</i>	<i>Char Name</i>	<i>Character Name</i>
'		close quote	fi	\(fi	fi
`		open quote	fl	\(fl	fl
-	\(em	3/4 Em dash	ff	\(ff	ff
-	-	hyphen or	ffi	\(Fi	ffi
-	\(hy	hyphen	ffl	\(Fl	ffl
-	\(-	current font minus	°	\(de	degree
•	\(bu	bullet	†	\(dg	dagger
□	\(sq	square	'	\(fm	foot mark
-	\(ru	rule	¢	\(ct	cent sign
¼	\(14	1/4	®	\(rg	registered
½	\(12	1/2	©	\(co	copyright
¾	\(34	3/4			

Non-ASCII characters and ' , ` , _ , + , - , = , and * on the special font.

The ASCII characters @, #, ", ;, <, >, \, {, }, ~, ^, and _ exist *only* on the special font and are printed as a 1-em space if that font is not mounted. The following characters exist only on the special font except for the upper case Greek letter names followed by † which are mapped into upper case English letters in whatever font is mounted on font position one (default Times Roman). The special math plus, minus, and equals are provided to insulate the appearance of equations from the choice of standard fonts.

<i>Input Character</i>	<i>Char Name</i>	<i>Character Name</i>	<i>Input Character</i>	<i>Char Name</i>	<i>Character Name</i>
+	\(pl	math plus	κ	\(*k	kappa
-	\(mi	math minus	λ	\(*l	lambda
=	\(eq	math equals	μ	\(*m	mu
*	\(**	math star	ν	\(*n	nu
§	\(sc	section	ξ	\(*c	xi
'	\(aa	acute accent	ο	\(*o	omicron
`	\(ga	grave accent	π	\(*p	pi
-	\(ul	underrule	ρ	\(*r	rho
/	\(sl	slash (matching backslash)	σ	\(*s	sigma
α	\(*a	alpha	ς	\(ts	terminal sigma
β	\(*b	beta	τ	\(*t	tau
γ	\(*g	gamma	υ	\(*u	upsilon
δ	\(*d	delta	φ	\(*f	phi
ε	\(*e	epsilon	χ	\(*x	chi
ζ	\(*z	zeta	ψ	\(*q	psi
η	\(*y	eta	ω	\(*w	omega
θ	\(*h	theta	Α	\(*A	Alpha†
ι	\(*i	iota	Β	\(*B	Beta†

<i>Char</i>	<i>Input Name</i>	<i>Character Name</i>
Γ	<code>\(*G</code>	Gamma
Δ	<code>\(*D</code>	Delta
E	<code>\(*E</code>	Epsilon†
Z	<code>\(*Z</code>	Zeta†
H	<code>\(*Y</code>	Eta†
Θ	<code>\(*H</code>	Theta
I	<code>\(*I</code>	Iota†
K	<code>\(*K</code>	Kappa†
Λ	<code>\(*L</code>	Lambda
M	<code>\(*M</code>	Mu†
N	<code>\(*N</code>	Nu†
Ξ	<code>\(*C</code>	Xi
O	<code>\(*O</code>	Omicron†
Π	<code>\(*P</code>	Pi
P	<code>\(*R</code>	Rho†
Σ	<code>\(*S</code>	Sigma
T	<code>\(*T</code>	Tau†
Y	<code>\(*U</code>	Upsilon
Φ	<code>\(*F</code>	Phi
X	<code>\(*X</code>	Chi†
Ψ	<code>\(*Q</code>	Psi
Ω	<code>\(*W</code>	Omega
$\sqrt{\quad}$	<code>\(sr</code>	square root
$\sqrt[n]{\quad}$	<code>\(rn</code>	root en extender
\geq	<code>\(>=</code>	$>=$
\leq	<code>\(<=</code>	$<=$
\equiv	<code>\(==</code>	identically equal
\approx	<code>\(≈</code>	approx =
\sim	<code>\(ap</code>	approximates
\neq	<code>\(!=</code>	not equal
\rightarrow	<code>\(-></code>	right arrow
\leftarrow	<code>\(<-</code>	left arrow
\uparrow	<code>\(ua</code>	up arrow
\downarrow	<code>\(da</code>	down arrow
\times	<code>\(mu</code>	multiply
\div	<code>\(di</code>	divide
\pm	<code>\(+-</code>	plus-minus
U	<code>\(cu</code>	cup (union)
\cap	<code>\(ca</code>	cap (intersection)
\subset	<code>\(sb</code>	subset of
\supset	<code>\(sp</code>	superset of
\subseteq	<code>\(ib</code>	improper subset
\supseteq	<code>\(ip</code>	improper superset
∞	<code>\(if</code>	infinity
∂	<code>\(pd</code>	partial derivative
∇	<code>\(gr</code>	gradient
\neg	<code>\(no</code>	not
\int	<code>\(is</code>	integral sign
\propto	<code>\(pt</code>	proportional to
\emptyset	<code>\(es</code>	empty set
\in	<code>\(mo</code>	member of

<i>Char</i>	<i>Input Name</i>	<i>Character Name</i>
	<code>\(br</code>	box vertical rule
‡	<code>\(dd</code>	double dagger
☞	<code>\(rh</code>	right hand
☜	<code>\(lh</code>	left hand
Ⓚ	<code>\(bs</code>	Bell System logo
	<code>\(or</code>	or
○	<code>\(ci</code>	circle
{	<code>\(lt</code>	left top of big curly bracket
	<code>\(lb</code>	left bottom
}	<code>\(rt</code>	right top
	<code>\(rb</code>	right bot
{	<code>\(lk</code>	left center of big curly bracket
}	<code>\(rk</code>	right center of big curly bracket
	<code>\(bv</code>	bold vertical
	<code>\(lf</code>	left floor (left bottom of big square bracket)
	<code>\(rf</code>	right floor (right bottom)
	<code>\(lc</code>	left ceiling (left top)
	<code>\(rc</code>	right ceiling (right top)

Addendum to the
NROFF/TROFF User's Manual
May 1977

Options

- h (NROFF only) Use output tabs during horizontal spacing to speed up output as well as to reduce output byte count. Device tab settings are assumed to be every 8 nominal character widths. The default settings of input (logical) tabs is also initialized to every 8 nominal character widths.
- z Efficiently suppresses formatted output. Only message output will occur (from **tm** requests and diagnostics).

Old Requests


- .ad *c* The adjustment type indicator *c* may now also be a number obtained from the “.j” register (see below).
- .so *name* The contents of file *name* will be interpolated at the point the **so** request is encountered. Previously, the interpolation was done upon return to the file-reading input level.


New Request

- .ab *text* Prints *text* on the message output and terminates without further processing. If *text* is missing, “User Abort.” is printed. Does *not* cause a break. The output buffer is flushed.

New Predefined Number Registers

- .k Read-only. Contains the horizontal size of the text portion (without indent) of the current partially-collected output line, if any, in the current environment.
- .j Read-only. Indicates the current adjustment mode and type. Can be saved and later given to the **ad** request to restore a previous mode.
- .P Read-only. Contains the value 1 if the current page is being printed, and is zero otherwise, i.e., if the current page did *not* appear in the **-o** option list.
- .L Read-only. Contains the current line-spacing parameter (the value of the most recent **ls** request).
- .c Provides general register access to the input line-number in the current input file. Contains the same value as the read-only “.c” register.







PWB/MM
Programmer's Workbench
Memorandum Macros

D. W. Smith
J. R. Mashey

October 1977

PWB/MM
Programmer's Workbench Memorandum Macros

CONTENTS

1. INTRODUCTION	1
1.1 Purpose	1
1.2 Conventions	1
1.3 Overall Structure of a Document	2
1.4 Definitions	2
1.5 Prerequisites and Further Reading	3
2. INVOKING THE MACROS	3
2.1 The mm Command	3
2.2 The -mm Flag	3
2.3 Typical Command Lines	4
2.4 Parameters that Can Be Set from the Command Line	5
2.5 Omission of -mm	6
3. FORMATTING CONCEPTS	6
3.1 Basic Terms	6
3.2 Arguments and Double Quotes	7
3.3 Unpaddable Spaces	7
3.4 Hyphenation	7
3.5 Tabs	8
3.6 Special Use of the BEL Character	8
3.7 Bullets	8
3.8 Dashes, Minus Signs, and Hyphens	8
3.9 Use of Formatter Requests	9
4. PARAGRAPHS AND HEADINGS	9
4.1 Paragraphs	9
4.2 Numbered Headings	9
4.3 Unnumbered Headings	12
4.4 Headings and the Table of Contents	12
4.5 First-Level Headings and the Page Numbering Style	12
4.6 User Exit Macros	13
4.7 Hints for Large Documents	14
5. LISTS	14
5.1 Basic Approach	14
5.2 Sample Nested Lists	14
5.3 Basic List Macros	15
5.4 List-Begin Macro and Customized Lists	18
6. MEMORANDUM AND RELEASED PAPER STYLES	19
6.1 Title	20
6.2 Author(s)	20
6.3 TM Number(s)	20
6.4 Abstract	21
6.5 Other Keywords	21
6.6 Memorandum Types	21
6.7 Date and Format Changes	22
6.8 Released-Paper Style	22
6.9 Order of Invocation of "Beginning" Macros	22
6.10 Example	23
6.11 Macros for the End of a Memorandum	23
6.12 Forcing a One-Page Letter	24

- 7. DISPLAYS 25
 - 7.1 Static Displays 25
 - 7.2 Floating Displays 25
 - 7.3 Tables 26
 - 7.4 Equations 26
 - 7.5 Figure, Table, and Equation Captions 26
 - 7.6 Blocks of Filled Text 27

- 8. FOOTNOTES 27
 - 8.1 Automatic Numbering of Footnotes 27
 - 8.2 Delimiting Footnote Text 27
 - 8.3 Format of Footnote Text • 28
 - 8.4 Spacing between Footnote Entries 29

- 9. PAGE HEADERS AND FOOTERS 29
 - 9.1 Default Headers and Footers 29
 - 9.2 Page Header 29
 - 9.3 Even-Page Header 29
 - 9.4 Odd-Page Header 30
 - 9.5 Page Footer 30
 - 9.6 Even-Page Footer 30
 - 9.7 Odd-Page Footer 30
 - 9.8 Footer on the First Page 30
 - 9.9 Default Header and Footer with “Section-Page” Numbering 30
 - 9.10 Use of Strings and Registers in Header and Footer Macros • 30
 - 9.11 Header and Footer Example • 31
 - 9.12 Generalized Top-of-Page Processing • 31
 - 9.13 Generalized Bottom-of-Page Processing 31

- 10. TABLE OF CONTENTS AND COVER SHEET 31
 - 10.1 Table of Contents 32
 - 10.2 Cover Sheet 33

- 11. MISCELLANEOUS FEATURES 33
 - 11.1 Bold, Italic, and Roman 33
 - 11.2 Justification of Right Margin 33
 - 11.3 SCCS Release Identification 34
 - 11.4 Two-Column Output 34
 - 11.5 Column Headings for Two-Column Output • 34
 - 11.6 Vertical Spacing 34
 - 11.7 Skipping Pages 35
 - 11.8 Setting Point Size and Vertical Spacing 35

- 12. ERRORS AND DEBUGGING 35
 - 12.1 Error Terminations 35
 - 12.2 Disappearance of Output 36

- 13. EXTENDING AND MODIFYING THE MACROS • 36
 - 13.1 Naming Conventions 36
 - 13.2 Sample Extensions 37

- 14. CONCLUSION 38
 - References 39
 - Appendix A: DEFINITIONS OF LIST MACROS • 41
 - Appendix B: USER-DEFINED LIST STRUCTURES • 42
 - Appendix C: SAMPLE FOOTNOTES 44
 - Appendix D: SAMPLE LETTER 46
 - Appendix E: ERROR MESSAGES 50
 - Appendix F: SUMMARY OF MACROS, STRINGS, AND NUMBER REGISTERS 52



PWB/MM—Programmer's Workbench Memorandum Macros

D. W. Smith

Bell Laboratories
Piscataway, New Jersey 08854

J. R. Mashey

Bell Laboratories
Murray Hill, New Jersey 07974

1. INTRODUCTION

1.1 Purpose

This memorandum is the user's guide and reference manual for PWB/MM (or just **-mm**), a general-purpose package of text formatting macros for use with the UNIX* text formatters *nroff* [9] and *troff* [9]. The purpose of PWB/MM is to provide to the users of PWB/UNIX a unified, consistent, and flexible tool for producing many common types of documents. Although PWB/UNIX provides other macro packages for various *specialized* formats, PWB/MM has become the standard, general-purpose macro package for most documents.

PWB/MM can be used to produce:

- Letters.
- Memoranda for File.
- Technical Memoranda.
- Released Papers.
- Programmer's Notes and other internal documents.
- Manuals.
- Books.
- etc.

The uses of PWB/MM range from single-page letters to documents of several hundred pages in length, such as user guides, design proposals, etc.

1.2 Conventions

Each section of this memorandum explains a single facility of PWB/MM. In general, the earlier a section occurs, the more necessary it is for most users. Some of the later sections can be completely ignored if PWB/MM defaults are acceptable. Likewise, each section progresses from normal-case to special-case facilities. We recommend reading a section in detail only until there is enough information to obtain the desired format, then skimming the rest of it, because some details may be of use to just a few people.

Numbers enclosed in curly brackets ({}) refer to section numbers within this document. For example, this is {1.2}.

Sections that require knowledge of the formatters {1.4} have a bullet (•) at the end of the section heading.

In the synopses of macro calls, square brackets ([]) surrounding an argument indicate that it is optional. Ellipses (...) show that the preceding argument may appear more than once.

A reference of the form *name*(N) points to page *name* in section N of the *PWB/UNIX User's Manual* [1].

The examples of *output* in this manual are as produced by *troff*; *nroff* output would, of course, look somewhat different (Appendix D shows *both* the *nroff* and *troff* output for a simple letter). In those

* UNIX is a Trademark of Bell Laboratories.

cases in which the behavior of the two formatters is truly different, the *nroff* action is described first, with the *troff* action following in parentheses. For example:

The title is underlined (bold).

means that the title is underlined in *nroff* and bold in *troff*.

1.3 Overall Structure of a Document

The input for a document that is to be formatted with PWB/MM possesses four major segments, any of which may be omitted; if present, they *must* occur in the following order:

- *Parameter-setting*—This segment sets the general style and appearance of a document. The user can control page width, margin justification, numbering styles for headings and lists, page headers and footers {9}, and many other properties of the document. Also, the user can add macros or redefine existing ones. This segment can be omitted entirely if one is satisfied with default values; it produces no actual output, but only performs the setup for the rest of the document.
- *Beginning*—This segment includes those items that occur only once, at the beginning of a document. Data (such as author, title, date, case numbers, etc.) for memoranda for file, technical memoranda, and released papers are contained in this segment {6}.
- *Body*—This segment is the actual text of the document. It may be as small as a single paragraph, or as large as hundreds of pages. It may have a hierarchy of *headings* up to seven levels deep {4}. Headings are automatically numbered (if desired) and can be saved to generate the table of contents. Six additional levels of subordination are provided by a set of *list* macros for automatic numbering, alphabetic sequencing, and “marking” of list items {5}. The body may also contain various types of displays, tables, figures, and footnotes {7, 8}.
- *Ending*—This segment contains those items that occur once only, at the end of a document. Included here are signature(s) and lists of notations (e.g., “copy to” lists) {6.12}. Certain macros may be invoked here to print information that is wholly or partially derived from the rest of the document, such as the table of contents or the cover sheet for a memorandum or a released paper {10}.

The existence and size of these four segments varies widely among different document types. Although a specific item (such as date, case number(s), author name(s), etc.) may be printed in several different ways depending on the document type, there is a uniform way of typing it in.

1.4 Definitions

The term *formatter* refers to either of the text-formatting programs *nroff* and *troff*.

Requests are built-in commands recognized by the formatters. Although one seldom needs to use these requests directly {3.9}, this document contains references to some of them. Full details are given in [9]. For example, the request:

```
.sp
```

inserts a blank line in the output.

Macros are named collections of requests. Each macro is an abbreviation for a collection of requests that would otherwise require repetition. PWB/MM supplies many macros, and the user can define additional ones. Macros and requests share the same set of names and are used in the same way.

Strings provide character variables, each of which names a string of characters. Strings are often used in page headers, page footers, and lists. They share the pool of names used by *requests* and *macros*. A string can be given a value via the *.ds* (define string) request, and its value can be obtained by referencing its name, preceded by “*” (for 1-character names) or “*(” (for 2-character names). For instance, the string *DT* in PWB/MM normally contains the current date, so that the *input* line:

```
Today is \*(DT.
```

may result in the following *output*:

```
Today is October 31, 1977.
```

The current date can be replaced, e.g.:

```
.ds DT 01/01/76
```

or by invoking a macro designed for that purpose {6.7.1}.

Number registers fill the role of integer variables. They are used for flags, for arithmetic, and for automatic numbering. A register can be given a value using a `.nr` request, and be referenced by preceding its name by “\n” (for 1-character names) or “\n(” (for 2-character names). For example, the following sets the value of the register *d* to 1 more than that of the register *dd*:

```
.nr d 1+\n(dd
```

See {13.1} regarding naming conventions for requests, macros, strings, and number registers.

1.5 Prerequisites and Further Reading

1.5.1 Prerequisites. We assume familiarity with UNIX at the level given in [3] and [4]. Some familiarity with the request summary in [9] is helpful.

1.5.2 Further Reading. [9] provides detailed descriptions of formatter capabilities, while [5] provides a general overview. See [6] (and possibly [7]) for instructions on formatting mathematical expressions. See *tbl(1)* and [11] for instructions on formatting tabular data.

Examples of formatted documents and of their respective input, as well as a quick reference to the material in this manual are given in [8].

2. INVOKING THE MACROS

This section tells how to access PWB/MM, shows PWB/UNIX command lines appropriate for various output devices, and describes command-line flags for PWB/MM. Note that file names, program names, and typical command sequences apply only to PWB/UNIX; different names and command lines may have to be used on other systems.

2.1 The `mm` Command

The `mm(1)` command can be used to print documents using *nroff* and PWB/MM; this command invokes *nroff* with the `-mm` flag {2.2}. It has options to specify preprocessing by *tbl(1)* and/or by *neqn(1)*, and for postprocessing by various output filters. Any arguments or flags that are not recognized by `mm(1)`, e.g. `-rC3`, are passed to *nroff* or to PWB/MM, as appropriate. The options, which can occur in any order but *must* appear before the file names, are:

- e *neqn(1)* is to be invoked.
- t *tbl(1)* is to be invoked.
- c *col(1)* is to be invoked.
- 12 need 12-pitch mode. Be sure that the pitch switch on the terminal is set to 12.
- 300 output is to a DASI300 terminal. This is the *default* terminal type.
- hp output is to a HP264x.
- 450 output is to a DASI450.
- 300S output is to a DASI300S.
- 300s output is to a DASI300S.
- tn output is to a GE TermiNet 300.
- tn300 output is to a GE TermiNet 300.
- ti output is to a Texas Instrument 700 series terminal.
- 37 output is to a TELETYPE[®] Model 37.

2.2 The `-mm` Flag

The PWB/MM package can also be invoked by including the `-mm` flag as an argument to the formatter. It causes the file `/usr/lib/tmac.m` to be read and processed before any other files. This action defines the PWB/MM macros, sets default values for various parameters, and initializes the formatter to be ready to process the files of input text.

2.3 Typical Command Lines

The prototype command lines are as follows (with the various options explained in {2.4} and in [9]).

- **Text** without tables or equations:

```
mm [options] filename ...  
or nroff [options] -mm filename ...  
or troff [options] -mm filename ...
```

- Text with tables:

```
mm -t [options] filename ...  
or tbl filename ... | nroff [options] -mm -  
or tbl filename ... | troff [options] -mm -
```

- Text with equations:

```
mm -e [options] filename ...  
or neqn filename ... | nroff [options] -mm -  
or eqn filename ... | troff [options] -mm -
```

- Text with both tables and equations:

```
mm -t -e [options] filename ...  
or tbl filename ... | neqn | nroff [options] -mm -  
or tbl filename ... | eqn | troff [options] -mm -
```

When formatting a document with *nroff*, the output should normally be processed for a specific type of terminal, because the output may require some features that are specific to a given terminal, e.g., reverse paper motion or half-line paper motion in both directions. Some commonly-used terminal types and the command lines appropriate for them are given below. See {2.4} as well as *gsi*(I), *hp*(I), *col*(I), and *terminals*(VII) for further information.

- DASI300 (GSI300/DTC300) in 10-pitch, 6 lines/inch mode and a line length of 65 characters:

```
mm filename ...  
or nroff -T300 -h -mm filename ...
```

- DASI300 (GSI300/DTC300) in 12-pitch, 6 lines/inch mode and a line length of 80—rather than 65—characters:

```
mm -12 filename ...  
or nroff -T300-12 -rW80 -rO3 -h -mm filename ...
```

or, equivalently (and more succinctly):

```
nroff -T300-12 -rT1 -h -mm filename ...
```

- DASI450 in 10-pitch, 6 lines/inch mode:

```
mm -450 filename ...  
or nroff -T450 -h -mm filename ...
```

- DASI450 in 12-pitch, 6 lines/inch mode:

```
mm -450 -12 filename ...  
or nroff -T450-12 -rW80 -rO3 -h -mm filename ...  
or nroff -T450-12 -rT1 -h -mm filename ...
```

- Hewlett-Packard HP264x CRT family:

```
mm -hp filename ...  
or nroff -h -mm filename ... | hp
```

- Any terminal incapable of reverse paper motion (GE TermiNet, Texas Instruments 700 series, etc.):

```
mm -tn filename ...  
or nroff -mm filename ... | col
```

- Versatec printer (see *vp(I)* for additional details):

```

vp [vp-options] "mm -rT2 -c filename ..."
or vp [vp-options] "nroff -rT2 -mm filename ... | col"

```

Of course, *tbl(I)* and *eqn(I)/neqn(I)*, if needed, must be invoked as shown in the command line prototypes at the beginning of this section.

If two-column processing {11.4} is used with *nroff*, the *-c* option must be specified to *mm(I)*, or the *nroff* output postprocessed by *col(I)*. In the latter case, the *-T37* terminal type must be specified to *nroff*, the *-h* option must *not* be specified, and the output of *col(I)* must be processed by the appropriate terminal filter (e.g., *gsi(I)*); *mm(I)* with the *-c* option handles all this automatically.

2.4 Parameters that Can Be Set from the Command Line

Number registers are commonly used within PWB/MM to hold parameter values that control various aspects of output style. Many of these can be changed within the text files via *.nr* requests. In addition, some of these registers can be set from the command line itself, a useful feature for those parameters that should *not* be permanently embedded within the input text itself. If used, these registers (with the possible exception of the register *P*—see below) *must* be set on the command line (or before the PWB/MM macro definitions are processed) and their meanings are:

- rA1 has the effect of invoking the *.AF* macro without an argument {6.7.2}.
- rB*n* defines the macros for the cover sheet and the table of contents. If *n* is 1, table-of-contents processing is enabled. If *n* is 2, then cover-sheet processing will occur. If *n* is 3, both will occur. That is, *B* having a value greater than 0 *defines* the *.TC* {10.1} and/or *.CS* {10.2} macros. Note that to have any effect, these macros must also be *invoked*.
- rC*n* *n* sets the type of copy (e.g., DRAFT) to be printed at the bottom of each page. See {9.5}.
n = 1 for OFFICIAL FILE COPY.
n = 2 for DATE FILE COPY.
n = 3 for DRAFT.
- rD1 sets *debug mode*. This flag requests the formatter to attempt to continue processing even if PWB/MM detects errors that would otherwise cause termination. It also includes some debugging information in the default page header {9.2, 11.3}.
- rL*k* sets the length of the physical page to *k* lines.¹ The default value is 66 lines per page. This parameter is used for obtaining 8 lines-per-inch output on 12-pitch terminals, or when directing output to a Versatec printer.
- rN*n* specifies the page numbering style. When *n* is 0 (default), all pages get the (prevailing) header {9.2}. When *n* is 1, the page header replaces the footer on page 1 only. When *n* is 2, the page header is omitted from page 1. When *n* is 3, "section-page" numbering {4.5} occurs.

<i>n</i>	<i>Page 1</i>	<i>Pages 2 ff.</i>
0	header	header
1	header replaces footer	header
2	no header	header
3	"section-page" as <i>footer</i>	

The contents of the prevailing header and footer do *not* depend on of the value of the number register *N*; *N* only controls whether and where the header (and, for *N*=3, the footer) is printed, as well as the page numbering style. In particular, if the header and footer are null {9.2, 9.5}, the value of *N* is irrelevant.

- rO*k* offsets output *k* spaces to the right.¹ It is helpful for adjusting output positioning on some terminals. NOTE: The register name is the capital letter "o", *not* the digit zero (0).
- rP*n* specifies that the pages of the document are to be numbered starting with *n*. This register may also be set via a *.nr* request in the input text.

¹ For *nroff*, *k* is an *unscaled* number representing lines or character positions; for *troff*, *k* must be *scaled*.

- rSn sets the point size and vertical spacing for the document. The default *n* is 10, i.e., 10-point type on 12-point leading (vertical spacing), giving 6 lines per inch {11.8}. This parameter applies to *nroff* only.
- rT*n* provides register settings for certain devices. If *n* is 1, then the line length and page offset are set for output directed to a DASI300 or DASI450 in 12-pitch, 6 lines/inch mode, i.e., they are set to 80 and 3, respectively. Setting *n* to 2 changes the page length to 84 lines per page and inhibits underlining; it is meant for output sent to the Versatec printer. The default value for *n* is 0. This parameter applies to *nroff* only.
- rU1 controls underlining of section headings. This flag causes only letters and digits to be underlined. Otherwise, all characters (including spaces) are underlined {4.2.2.4.2}. This parameter applies to *nroff* only.
- rW*k* page width (i.e., line length and title length) is set to *k*.² This can be used to change the page width from the default value of 65 characters (6.5 inches).

2.5 Omission of -mm

If a large number of arguments are required on the command line, it may be convenient to set up the first (or only) input file of a document as follows:

```
zero or more initializations of registers listed in {2.4}
.so /usr/lib/tmac.m
remainder of text
```

In this case, one must *not* use the **-mm** flag (nor the *mm(I)* command); the *.so* request has the equivalent effect, but the registers in {2.4} must be initialized *before* the *.so* request, because their values are meaningful only if set before the macro definitions are processed. When using this method, it is best to “lock” into the input file only those parameters that are seldom changed. For example:

```
.nr W 80
.nr O 10
.nr N 3
.nr B 1
.so /usr/lib/tmac.m
.H 1 "INTRODUCTION"
:
```

specifies, for *nroff*, a line length of 80, a page offset of 10, “section-page” numbering, and table of contents processing.

3. FORMATTING CONCEPTS

3.1 Basic Terms

The normal action of the formatters is to *fill* output lines from one or more input lines. The output lines may be *justified* so that both the left and right margins are aligned. As the lines are being filled, words are hyphenated {3.4} as necessary. It is possible to turn any of these modes on and off (see *.SA* {11.2}, *Hy* {3.4}, and the formatter *.nf* and *.fi* requests [9]). Turning off fill mode also turns off justification and hyphenation.

Certain formatting commands (requests and macros) cause the filling of the current output line to cease, the line (of whatever length) to be printed, and the subsequent text to begin a new output line. This printing of a partially filled output line is known as a *break*. A few formatter requests and most of the PWB/MM macros cause a break.

While formatter requests can be used with PWB/MM, one must fully understand the consequences and side-effects that each such request might have. Actually, there is little need to use formatter requests; the macros described here should be used in most cases because:

2. For *nroff*, *k* is an *unscaled* number representing lines or character positions; for *troff*, *k* must be *scaled*.

- it is much easier to control (and change at any later point in time) the overall style of the document.
- complicated facilities (such as footnotes or tables of contents) can be obtained with ease.
- the user is insulated from the peculiarities of the formatter language.

A good rule is to use formatter requests only when absolutely necessary {3.9}.

In order to make it easy to revise the input text at a later time, input lines should be kept short and should be broken at the end of clauses; each new full *sentence must* begin on a new line.

3.2 Arguments and Double Quotes

For any macro call, a *null argument* is an argument whose width is zero. Such an argument often has a special meaning; the preferred form for a null argument is "". Note that *omitting* an argument is *not* the same as supplying a *null argument* (for example, see the .MT macro in {6.6}). Furthermore, omitted arguments can occur only at the end of an argument list, while null arguments can occur anywhere.

Any macro argument containing ordinary (paddable) spaces *must* be enclosed in double quotes (").³ Otherwise, it will be treated as several separate arguments.

Double quotes (") are *not* permitted as part of the value of a macro argument or of a string that is to be used as a macro argument. If you must, use two grave accents (`) and/or two acute accents (´) instead. This restriction is necessary because many macro arguments are processed (interpreted) a variable number of times; for example, headings are first printed in the text and may be (re)printed in the table of contents.

3.3 Unpaddable Spaces

When output lines are *justified* to give an even right margin, existing spaces in a line may have additional spaces appended to them. This may harm the desired alignment of text. To avoid this problem, it is necessary to be able to specify a space that cannot be expanded during justification, i.e., an *unpaddable space*. There are several ways to accomplish this.

First, one may type a backslash followed by a space (" \ "). This pair of characters directly generates an *unpaddable space*. Second, one may sacrifice some seldom-used character to be translated into a space upon output. Because this translation occurs after justification, the chosen character may be used anywhere an unpaddable space is desired. The tilde (~) is often used for this purpose. To use it in this way, insert the following at the beginning of the document:

```
.tr ~
```

If a tilde must actually appear in the output, it can be temporarily "recovered" by inserting:

```
.tr ^^
```

before the place where it is needed. Its previous usage is restored by repeating the ".tr ~", but only after a break or after the line containing the tilde has been forced out. Note that the use of the tilde in this fashion is *not* recommended for documents in which the tilde is used within equations.

3.4 Hyphenation

The formatters (and, therefore, PWB/MM) will automatically hyphenate words, if need be. However, the user may specify the hyphenation points for a specific occurrence of any word by the use of a special character known as a hyphenation indicator, or may specify hyphenation points for a small list of words (about 128 characters).

If the *hyphenation indicator* (initially, the two-character sequence "\%") appears at the beginning or end of a word, the word is *not* hyphenated. Alternatively, it can be used to indicate legal hyphenation point(s) inside a word. In any case, *all* occurrences of the hyphenation indicator disappear on output.

The user may specify a different hyphenation indicator:

```
.HC [hyphenation-indicator]
```

3. A double quote (") is a *single* character that must not be confused with two apostrophes or acute accents (´), or with two grave accents (`).

The circumflex (^) is often used for this purpose; this is done by inserting the following at the beginning of a document:

```
.HC ^
```

Note that any word containing hyphens or dashes—also known as *em* dashes—will be hyphenated immediately after a hyphen or dash if it is necessary to hyphenate the word, *even if the formatter hyphenation function is turned off*.

Hyphenation can be turned off in the body of the text by specifying:

```
.nr Hy 0
```

once at the beginning of the document. For hyphenation control within footnote text and across pages, see {8.3}.

The user may supply, via the .hw request, a small list of words with the proper hyphenation points indicated. For example, to indicate the proper hyphenation of the word “printout,” one may specify:

```
.hw print-out
```

3.5 Tabs

The macros .MT {6.6}, .TC {10.1}, and .CS {10.2} use the formatter .ta request to set tab stops, and then restore the *default* values⁴ of tab settings. Thus, setting tabs to other than the default values is the user’s responsibility.

Note that a tab character is always interpreted with respect to its position on the *input line*, rather than its position on the output line. In general, tab characters should appear only on lines processed in “no-fill” mode {3.1}.

Also note that *tbl(I)* {7.3} changes tab stops, but does *not* restore the default tab settings.

3.6 Special Use of the BEL Character

The non-printing character BEL is used as a delimiter in many macros where it is necessary to compute the width of an argument or to delimit arbitrary text, e.g., in headers and footers {9}, headings {4}, and list marks {5}. Users who include BEL characters in their input text (especially in arguments to macros) will receive mangled output.

3.7 Bullets

A bullet (•) is often obtained on a typewriter terminal by using an “o” overstruck by a “+”. For compatibility with *troff*, a bullet string is provided by PWB/MM. Rather than overstriking, use the sequence:

```
\*(BU
```

wherever a bullet is desired. Note that the bullet list (.BL) macros {5.3.3.2} use this string to automatically generate the bullets for the list items.

3.8 Dashes, Minus Signs, and Hyphens

Troff has distinct graphics for a dash, a minus sign, and a hyphen, while *nroff* does not. Those who intend to use *nroff* only may use the minus sign (“-”) for all three.

Those who wish mainly to use *troff* should follow the escape conventions of {9}.

Those who want to use both formatters must take care during text preparation. Unfortunately, these characters cannot be represented in a way that is both compatible and convenient. We suggest the following approach:

Dash Type “ -- ” for each text dash. These can be left alone for *nroff*, and later globally translated for *troff* to “\em”, namely an *em* dash (—). Note that the dash list (.DL) macros {5.3.3.3} automatically generate the *em* dashes for the list items.

4. Every eight characters in *nroff*; every 1/2 inch in *troff*.

Hyphen Type “-” and use as is for both formatters. *Nroff* will print it as is, and *troff* will print a true hyphen.

Minus Type “\” for a true minus sign, regardless of formatter. *Nroff* will effectively ignore the “\”, while *troff* will print a true minus sign.

3.9 Use of Formatter Requests

Most formatter requests [9] should *not* be used with PWB/MM because PWB/MM provides the corresponding formatting functions in a much more user-oriented and surprise-free fashion than do the basic formatter requests [3.1]. However, some formatter requests *are* useful with PWB/MM, namely:

.af	.br	.ce	.de	.ds	.fi	.hw	.ls	.nf	.nr
.nx	.rm	.rr	.rs	.so	.sp	.ta	.ti	.tl	.tr

The .fp, .lg, and .ss requests are also sometimes useful for *troff*. Use of other requests without fully understanding their implications very often leads to disaster.

4. PARAGRAPHS AND HEADINGS

This section describes simple paragraphs and section headings. Additional paragraph and list styles are covered in {5}.

4.1 Paragraphs

.P [type]
one or more lines of text.

This macro is used to begin two kinds of paragraphs. In a *left-justified* paragraph, the first line begins at the left margin, while in an *indented* paragraph, it is indented five spaces (see below).

A document possesses a *default paragraph style* obtained by specifying “.P” before each paragraph that does *not* follow a heading [4.2]. The default style is controlled by the register *Pt*. The initial value of *Pt* is 2, which provides indented paragraphs *except* after headings, lists, and displays, in which case they are left-justified. All paragraphs can be forced to be left-justified by inserting the following at the beginning of the document:

.nr Pt 0

All paragraphs can be forced to be indented by inserting:

.nr Pt 1

at the beginning of the document.

The amount a paragraph is indented is contained in the register *Pi*, whose default value is 5. To indent paragraphs by, say, 10 spaces, insert:

.nr Pi 10

at the beginning of the document. Of course, both the *Pi* and *Pt* register values must be greater than zero for any paragraphs to be indented.

▀ *Values that specify indentation must be unscaled and are treated as “character positions,” i.e., as a number of ens. In troff, an en is the number of points (1 point = 1/72 of an inch) equal to half the current point size. In nroff, an en is equal to the width of a character.*

Regardless of the value of *Pt*, an *individual* paragraph can be forced to be left-justified or indented. “.P 0” always forces left justification; “.P 1” always causes indentation by the amount specified by the register *Pi*.

If .P occurs inside a *list*, the indent (if any) of the paragraph is added to the current list indent {5}.

4.2 Numbered Headings

.H level [heading-text]
zero or more lines of text

The .H macro provides seven levels of numbered headings, as illustrated by this document. Level 1 is the most major or highest; level 7 the lowest.

☛ *There is no need for a .P macro after a .H (or .HU {4.3}), because the .H macro also performs the function of the .P macro. In fact, if a .P follows a .H, the user loses much of the flexibility provided by the .H mechanism {4.2.2.2}.*

4.2.1 Normal Appearance. The normal appearance of headings is as shown in this document. The effect of .H varies according to the *level* argument. First-level headings are *preceded* by two blank lines (one vertical space); all others are *preceded* by one blank line (½ a vertical space).

.H 1 heading-text gives an underlined (bold) heading *followed* by a single blank line (½ a vertical space). The following text begins on a new line and is indented according to the current paragraph type. Full capital letters should normally be used to make the heading stand out.

.H 2 heading-text yields an underlined (bold) heading followed by a single blank line (½ a vertical space). The following text begins on a new line and is indented according to the current paragraph type. Normally, initial capitals are used.

.H *n* heading-text for $3 \leq n \leq 7$, produces an underlined (italic) heading followed by two spaces. The following text appears on the same line, i.e., these are *run-in* headings.

Appropriate numbering and spacing (horizontal and vertical) occur even if the heading text is omitted from a .H macro call.

Here are the first few .H calls of {4}:

```
.H 1 "PARAGRAPHS AND HEADINGS"  
.H 2 "Paragraphs"  
.H 2 "Numbered Headings"  
.H 3 "Normal Appearance."  
.H 3 "Altering Appearance of Headings."  
.H 4 "Pre-Spacing and Page Ejection."  
.H 4 "Spacing After Headings."  
.H 4 "Centered Headings."  
.H 4 "Bold, Italic, and Underlined Headings."  
.H 5 "Control by Level."
```

4.2.2 Altering Appearance of Headings. Users satisfied with the default appearance of headings may skip to {4.3}. One can modify the appearance of headings quite easily by setting certain registers and strings at the beginning of the document. This permits quick alteration of a document's style, because this style-control information is concentrated in a few lines, rather than being distributed throughout the document.

4.2.2.1 Pre-Spacing and Page Ejection. A first-level heading normally has two blank lines (one vertical space) preceding it, and all others have one blank line (½ a vertical space). If a multi-line heading were to be split across pages, it is automatically moved to the top of the next page. Every first-level heading may be forced to the top of a new page by inserting:

```
.nr Ej 1
```

at the beginning of the document. Long documents may be made more manageable if each section starts on a new page. Setting *Ej* to a higher value causes the same effect for headings up to that level, i.e., a page eject occurs if the heading level is less than or equal to *Ej*.

4.2.2.2 Spacing After Headings. Three registers control the appearance of text immediately following a .H call. They are *Hb* (heading break level), *Hs* (heading space level), and *Hi* (post-heading indent).

If the heading level is less than or equal to *Hb*, a break {3.1} occurs after the heading. If the heading level is less than or equal to *Hs*, a blank line (½ a vertical space) is inserted after the heading. Defaults for *Hb* and *Hs* are 2. If a heading level is greater than *Hb* and also greater than *Hs*, then the

heading (if any) is run into the following text. These registers permit headings to be separated from the text in a consistent way throughout a document, while allowing easy alteration of white space and heading emphasis.

For any *stand-alone* heading, i.e., a heading not run into the following text, the alignment of the next line of output is controlled by the register *Hi*. If *Hi* is 0, text is left-justified. If *Hi* is 1 (the *default* value), the text is indented according to the paragraph type as specified by the register *Pr* {4.1}. Finally, if *Hi* is 2, text is indented to line up with the first word of the heading itself, so that the heading number stands out more clearly. Note that this feature is defeated if a .P macro follows the .H or .HU macro {4.2}.

For example, to cause a blank line (1/2 a vertical space) to appear after the first three heading levels, to have no run-in headings, and to force the text following all headings to be left-justified (regardless of the value of *Pr*), the following should appear at the top of the document:

```
.nr Hs 3
.nr Hb 7
.nr Hi 0
```

4.2.2.3 *Centered Headings.* The register *Hc* can be used to obtain centered headings. A heading is centered if its level is less than or equal to *Hc*, and if it is also stand-alone {4.2.2.2}. *Hc* is 0 initially (no centered headings).

4.2.2.4 *Bold, Italic, and Underlined Headings.*

4.2.2.4.1 *Control by Level.* Any heading that is underlined by *nroff* is made bold or italic by *troff*. The string *HF* (heading font) contains seven codes that specify the fonts for heading levels 1-7. The legal codes, their interpretations, and the defaults for *HF* are:

Formatter	HF Code			Default HF
	1	2	3	
nroff	no underline	underline	underline	3 3 2 2 2 2 2
troff	roman	italic	bold	3 3 2 2 2 2 2

Thus, all levels are underlined in *nroff*; in *troff*, levels 1 and 2 are bold, levels 3 through 7 are italic. The user may reset *HF* as desired. Any value omitted from the right end of the list is taken to be 1. For example, the following would result in five underlined (bold) levels and two non-underlined (roman) levels:

```
.ds HF 3 3 3 3 3
```

4.2.2.4.2 *Nroff Underlining Style.* *Nroff* can underline in two ways. The normal style (.ul request) is to underline only letters and digits. The continuous style (.cu request) underlines all characters, including spaces. By default, PWB/MM attempts to use the continuous style on any heading that is to be underlined, is *not* run-in, and is short enough to fit on a single line. If a heading is to be underlined, but is either run-in or is too long, it is underlined the normal way (i.e., only letters and digits are underlined).

All underlining of headings can be forced to the normal way by using the -rU1 flag when invoking *nroff* {2.4}.

4.2.2.5 *Marking Styles—Numerals and Concatenation.*

```
.HM [arg1] ... [arg7]
```

The registers named *H1* through *H7* are used as counters for the seven levels of headings. Their values are normally printed using Arabic numerals. The .HM macro (heading mark style) allows this choice to be overridden, thus providing "outline" and other document styles. This macro can have up to seven arguments; each argument is a string indicating the type of marking to be used. Legal values and their meanings are shown below; omitted values are interpreted as 1, while illegal values have no effect.

<i>Value</i>	<i>Interpretation</i>
1	Arabic (default for all levels)
0001	Arabic with enough leading zeroes to get the specified number of digits
A	Upper-case alphabetic
a	Lower-case alphabetic
I	Upper-case Roman
i	Lower-case Roman

By default, the complete heading mark for a given level is built by concatenating the mark for that level to the right of all marks for all levels of higher value. To inhibit the concatenation of heading level marks, i.e., to obtain just the current level mark followed by a period, set the register *Ht* (heading-mark type) to 1.

For example, a commonly-used “outline” style is obtained by:

```
.HM I A 1 a i  
.nr Ht 1
```

4.3 Unnumbered Headings

.HU heading-text

.HU is a special case of .H; it is handled in the same way as .H, except that no heading mark is printed. In order to preserve the hierarchical structure of headings when .H and .HU calls are intermixed, each .HU heading is considered to exist at the level given by register *Hu*, whose initial value is 2. Thus, in the normal case, the only difference between:

.HU heading-text

and

.H 2 heading-text

is the printing of the heading mark for the latter. Both have the effect of incrementing the numbering counter for level 2, and resetting to zero the counters for levels 3 through 7. Typically, the value of *Hu* should be set to make unnumbered headings (if any) be the lowest-level headings in a document.

.HU can be especially helpful in setting up Appendices and other sections that may not fit well into the numbering scheme of the main body of a document {13.2.1}.

4.4 Headings and the Table of Contents

The text of headings and their corresponding page numbers can be automatically collected for a table of contents. This is accomplished by doing the following three things:

- specifying in the register *Ci* what level headings are to be saved;
- invoking the .TC macro {10.1} at the end of the document;
- and specifying -rB*n* {2.4} on the command line.

Any heading whose level is less than or equal to the value of the register *Ci* (contents level) is saved and later displayed in the table of contents. The default value for *Ci* is 2, i.e., the first two levels of headings are saved.

Due to the way the headings are saved, it is possible to exceed the formatter’s storage capacity, particularly when saving many levels of many headings, while also processing displays {7} and footnotes {8}. If this happens, the “Out of temp file space” diagnostic {Appendix E} will be issued; the only remedy is to save fewer levels and/or to have fewer words in the heading text.

4.5 First-Level Headings and the Page Numbering Style

By default, pages are numbered sequentially at the top of the page. For large documents, it may be desirable to use page numbering of the form “section-page,” where *section* is the number of the current first-level heading. This page numbering style can be achieved by specifying the flag -rN3 on the command line {9.9}. As a side effect, this also has the effect of setting *Ej* to 1, i.e., each section

begins on a new page. In this style, the page number is printed at the *bottom* of the page, so that the correct section number is printed.

4.6 User Exit Macros •

▀ This section is intended only for users who are accustomed to writing formatter macros.

```
.HX dlevel rlevel heading-text
.HZ dlevel rlevel heading-text
```

The .HX and .HZ macros are the means by which the user obtains a final level of control over the previously-described heading mechanism. PWB/MM does not define .HX and .HZ; they are intended to be defined by the user. The .H macro invokes .HX shortly before the actual heading text is printed; it calls .HZ as its last action. All the default actions occur if these macros are not defined. If the .HX or .HZ (or both) are defined by the user, the user-supplied definition is interpreted at the appropriate point. These macros can therefore influence the handling of all headings, because the .HU macro is actually a special case of the .H macro.

If the user originally invoked the .H macro, then the derived level (*dlevel*) and the real level (*rlevel*) are both equal to the level given in the .H invocation. If the user originally invoked the .HU macro {4.5}, *dlevel* is equal to the contents of register *Hu*, and *rlevel* is 0. In both cases, *heading-text* is the text of the original invocation.

By the time .H calls .HX, it has already incremented the heading counter of the specified level {4.2.2.5}, produced blank line(s) (vertical space) to precede the heading {4.2.2.1}, and accumulated the “heading mark”, i.e., the string of digits, letters, and periods needed for a numbered heading. When .HX is called, all user-accessible registers and strings can be referenced, as well as the following:

- | | |
|-------------|--|
| string }0 | If <i>rlevel</i> is non-zero, this string contains the “heading mark.” Two unpadding spaces (to separate the <i>mark</i> from the <i>heading</i>) have been appended to this string. If <i>rlevel</i> is 0, this string is null. |
| register ;0 | This register indicates the type of spacing that is to follow the heading {4.2.2.2}. A value of 0 means that the heading is run-in. A value of 1 means a break (but no blank line) is to follow the heading. A value of 2 means that a blank line (½ a vertical space) is to follow the heading. |
| string }2 | If register ;0 is 0, this string contains two unpadding spaces that will be used to separate the (run-in) <i>heading</i> from the following <i>text</i> . If register ;0 is non-zero, this string is null. |
| register ;3 | This register contains an adjustment factor for a .ne request issued before the heading is actually printed. On entry to .HX, it has the value 3 if <i>dlevel</i> equals 1, and 1 otherwise. The .ne request is for the following number of lines: the contents of the register ;0 taken as blank lines (halves of vertical space) plus the contents of register ;3 as blank lines (halves of vertical space) plus the number of lines of the heading. |

The user may alter the values of }0, }2, and ;3 within .HX as desired. The following are examples of actions that might be performed by defining .HX to include the lines shown:

```
Change first-level heading mark from format n. to n.0:
.if \\$1=1 .ds }0 \\n(H1.0\□\□                               (□ stands for a space)

Separate run-in heading from the text with a period and two unpadding spaces:
.if \\n(;0=0 .ds }2 .\□\□

Assure that at least 15 lines are left on the page before printing a first-level heading:
.if \\$1=1 .nr ;3 15-\\n(;0

Add 3 additional blank lines before each first-level heading:
.if \\$1=1 .sp 3
```

If temporary string or macro names are used within .HX, care must be taken in the choice of their names {13.1}.

.HZ is called at the end of .H to permit user-controlled actions after the heading is produced. For example, in a large document, sections may correspond to chapters of a book, and the user may want to reset counters for footnotes, figures, tables, etc. Another use might be to change a page header or footer. For example:

```
.de HZ
.if \\$1=1 \\{.nr :p 0 \" footnotes
.   nr Fg 0 \" figures
.   nr Tb 0 \" tables
.   nr Ec 0 \" equations
.   PF \"Section \\$3\"\\}
..
```

4.7 Hints for Large Documents

A large document is often organized for convenience into one file per section. If the files are numbered, it is wise to use enough digits in the names of these files for the maximum number of sections, i.e., use suffix numbers 01 through 20 rather than 1 through 9 and 10 through 20.

Users often want to format individual sections of long documents. To do this with the correct section numbers, it is necessary to set register *HI* to 1 less than the number of the section just *before* the corresponding “.H 1” call. For example, at the beginning of section 5, insert:

```
.nr HI 4
```

■ *This is a dangerous practice: it defeats the automatic (re)numbering of sections when sections are added or deleted. Remove such lines as soon as possible.*

5. LISTS

This section describes many different kinds of lists: automatically-numbered and alphabetized lists, bullet lists, dash lists, lists with arbitrary marks, and lists starting with arbitrary strings, e.g., with terms or phrases to be defined.

5.1 Basic Approach

In order to avoid repetitive typing of arguments to describe the appearance of items in a list, PWB/MM provides a convenient way to specify lists. All lists are composed of the following parts:

- A *list-initialization* macro that controls the appearance of the list: line spacing, indentation, marking with special symbols, and numbering or alphabetizing.
- One or more *List Item* (.LI) macros, each followed by the actual text of the corresponding list item.
- The *List End* (.LE) macro that terminates the list and restores the previous indentation.

Lists may be nested up to six levels. The list-initialization macro saves the previous list status (indentation, marking style, etc.); the .LE macro restores it.

With this approach, the format of a list is specified only once at the beginning of that list. In addition, by building on the existing structure, users may create their own customized sets of list macros with relatively little effort {5.4, Appendix A, Appendix B}.

5.2 Sample Nested Lists

The input for several lists and the corresponding output are shown below. The .AL and .DL macro calls {5.3.3} contained therein are examples of the *list-initialization* macros. This example will help us to explain the material in the following sections. Input text:

.AL A

.LI

This is an alphabetized item.

This text shows the alignment of the second line of the item.

The quick brown fox jumped over the lazy dog's back.

.AL

.LI

This is a numbered item.

This text shows the alignment of the second line of the item.

The quick brown fox jumped over the lazy dog's back.

.DL

.LI

This is a dash item.

This text shows the alignment of the second line of the item.

The quick brown fox jumped over the lazy dog's back.

.LI + 1

This is a dash item with a "plus" as prefix.

This text shows the alignment of the second line of the item.

The quick brown fox jumped over the lazy dog's back.

.LE

.LI

This is numbered item 2.

.LE

.LI

This is another alphabetized item, B.

This text shows the alignment of the second line of the item.

The quick brown fox jumped over the lazy dog's back.

.LE

.P

This paragraph appears at the left margin.

Output:

- A. This is an alphabetized item. This text shows the alignment of the second line of the item. The quick brown fox jumped over the lazy dog's back.
 - 1. This is a numbered item. This text shows the alignment of the second line of the item. The quick brown fox jumped over the lazy dog's back.
 - This is a dash item. This text shows the alignment of the second line of the item. The quick brown fox jumped over the lazy dog's back.
 - + - This is a dash item with a "plus" as prefix. This text shows the alignment of the second line of the item. The quick brown fox jumped over the lazy dog's back.
 - 2. This is numbered item 2.
- B. This is another alphabetized item, B. This text shows the alignment of the second line of the item. The quick brown fox jumped over the lazy dog's back.

This paragraph appears at the left margin.

5.3 Basic List Macros

Because all lists share the same overall structure except for the list-initialization macro, we first discuss the macros common to all lists. Each list-initialization macro is covered in {5.3.3}.

5.3.1 List Item.

.LI [mark] [1]

one or more lines of text that make up the list item.

The `.LI` macro is used with all lists. It normally causes the output of a single blank line (½ a vertical space) before its item, although this may be suppressed. If no arguments are given, it labels its item with the *current mark*, which is specified by the most recent list-initialization macro. If a single argument is given to `.LI`, that argument is output *instead of* the current mark. If two arguments are given, the first argument becomes a *prefix* to the current mark, thus allowing the user to emphasize one or more items in a list. One unpaddingable space is inserted between the prefix and the mark. For example:

```
.BL 6
.LI
This is a simple bullet item.
.LI +
This replaces the bullet with a "plus."
.LI + 1
But this uses "plus" as prefix to the bullet.
.LE
```

yields:

- This is a simple bullet item.
- + This replaces the bullet with a "plus."
- + • But this uses "plus" as prefix to the bullet.

■ *The mark must not contain ordinary (paddingable) spaces, because alignment of items will be lost if the right margin is justified {3.3}.*

If the *current mark* (in the *current list*) is a null string, and the first argument of `.LI` is omitted or null, the resulting effect is that of a *hanging indent*, i.e., the first line of the following text is "outdented," starting at the same place where the *mark* would have started {5.3.3.6}.

5.3.2 List End.

```
.LE [1]
```

List End restores the state of the list back to that existing just before the most recent list-initialization macro call. If the optional argument is given, the `.LE` outputs a blank line (½ a vertical space). This option should generally be used only when the `.LE` is followed by running text, but not when followed by a macro that produces blank lines of its own, such as `.P`, `.H`, or `.LI`.

`.H` and `.HU` automatically clear all list information, so one may legally omit the `.LE(s)` that would normally occur just before either of these macros. Such a practice is *not* recommended, however, because errors will occur if the list text is separated from the heading at some later time (e.g., by insertion of text).

5.3.3 List Initialization Macros. The following are the various list-initialization macros. They are actually implemented as calls to the more basic `.LB` macro {5.4}.

5.3.3.1 Automatically-Numbered or Alphabetized Lists.

```
.AL [type] [text-indent] [1]
```

The `.AL` macro is used to begin sequentially-numbered or alphabetized lists. If there are no arguments, the list is numbered, and text is indented L_i (initially 5)⁵ spaces from the indent in force when the `.AL` is called, thus leaving room for two digits, a period, and two spaces before the text.

The *type* argument may be given to obtain a different type of sequencing, and its value should indicate the first element in the sequence desired, i.e., it must be 1, A, a, I, or i {4.2.2.5}.⁶ If *type* is omitted or null, then "1" is assumed. If *text-indent* is non-null, it is used as the number of spaces from the current indent to the text, i.e., it is used instead of L_i for this list only. If *text-indent* is null, then the value of L_i will be used.

5. Values that specify indentation must be *unscaled* and are treated as "character positions," i.e., as the number of *ens*.

6. Note that the "0001" format is *not* permitted.

If the third argument is given, a blank line (½ a vertical space) will *not* separate the items in the list. A blank line (½ a vertical space) will occur before the first item, however.

5.3.3.2 *Bullet List.*

.BL [text-indent] [1]

.BL begins a bullet list, in which each item is marked by a bullet (●) followed by one space. If *text-indent* is non-null, it overrides the default indentation—the amount of paragraph indentation as given in the register *Pi* {4.1}.⁷

If a second argument is specified, no blank lines will separate the items in the list.

5.3.3.3 *Dash List.*

.DL [text-indent] [1]

.DL is identical to **.BL**, except that a dash is used instead of a bullet.

5.3.3.4 *Marked List.*

.ML mark [text-indent] [1]

.ML is much like **.BL** and **.DL**, but expects the user to specify an arbitrary mark, which may consist of more than a single character. Text is indented *text-indent* spaces if the second argument is not null; otherwise, the text is indented one more space than the width of *mark*. If the third argument is specified, no blank lines will separate the items in the list.

■ *The mark must not contain ordinary (paddable) spaces, because alignment of items will be lost if the right margin is justified {3.3}.*

5.3.3.5 *Reference List.*

.RL [text-indent] [1]

A **.RL** call begins an automatically-numbered list in which the numbers are enclosed by square brackets ([]). *Text-indent* may be supplied, as for **.AL**. If omitted or null, it is assumed to be 6, a convenient value for lists numbered up to 99. If the second argument is specified, no blank lines will separate the items in the list. The list of references {14} was produced using the **.RL** macro.

5.3.3.6 *Variable-Item List.*

.VL text-indent [mark-indent] [1]

When a list begins with a **.VL**, there is effectively no *current mark*; it is expected that each **.LI** will provide its own mark. This form is typically used to display definitions of terms or phrases. *Mark-indent* gives the number of spaces from the current indent to the beginning of the *mark*, and it defaults to 0 if omitted or null. *Text-indent* gives the distance from the current indent to the beginning of the text. If the third argument is specified, no blank lines will separate the items in the list. Here is an example of **.VL** usage:

7. So that, in the default case, the text of bullet and dash lists lines up with the first line of indented paragraphs.

```
.tr ~
.VL 20 2
.LI mark~1
Here is a description of mark 1;
`mark 1` of the .LI line contains a tilde translated to an unpaddable space in order
to avoid extra spaces between
`mark` and `1` {3.3}.
.LI second~mark
This is the second mark, also using a tilde translated to an unpaddable space.
.LI third~mark~longer~than~indent:
This item shows the effect of a long mark; one space separates the mark
from the text.
.LI ~
This item effectively has no mark because the
tilde following the .LI is translated into a space.
.LE
```

yields:

```
mark 1      Here is a description of mark 1; "mark 1" of the .LI line contains a tilde
            translated to an unpaddable space in order to avoid extra spaces between "mark"
            and "1" {3.3}.

second mark  This is the second mark, also using a tilde translated to an unpaddable space.

third mark longer than indent: This item shows the effect of a long mark; one space separates the
            mark from the text.

            This item effectively has no mark because the tilde following the .LI is translated
            into a space.
```

The tilde argument on the last .LI above is required; otherwise a *hanging indent* would have been produced. A *hanging indent* is produced by using .VL and calling .LI with no arguments or with a null first argument. For example:

```
.VL 10
.LI
Here is some text to show a hanging indent.
The first line of text is at the left margin.
The second is indented 10 spaces.
.LE
```

yields:

Here is some text to show a hanging indent. The first line of text is at the left margin. The second is indented 10 spaces.

■ *The mark must not contain ordinary (paddable) spaces, because alignment of items will be lost if the right margin is justified {3.3}.*

5.4 List-Begin Macro and Customized Lists •

```
.LB text-indent mark-indent pad type [mark] [LI-space] [LB-space]
```

The list-initialization macros described above suffice for almost all cases. However, if necessary, one may obtain more control over the layout of lists by using the basic list-begin macro .LB, which is also used by all the other list-initialization macros (Appendix A). Its arguments are as follows:

Text-indent gives the number of spaces that the text is to be indented from the current indent. Normally, this value is taken from the register *Li* for automatic lists and from the register *Pi* for bullet and dash lists.

The combination of *mark-indent* and *pad* determines the placement of the mark. The mark is placed within an area (called *mark area*) that starts *mark-indent* spaces to the right of the current indent, and

ends where the text begins (i.e., ends *text-indent* spaces to the right of the current indent).⁸ Within the **mark** area, the mark is *left-justified* if *pad* is 0. If *pad* is greater than 0, say *n*, then *n* blanks are appended to the mark; the *mark-indent* value is ignored. The resulting string immediately precedes the text. That is, the mark is effectively *right-justified pad* spaces immediately to the left of the text.

Type and *mark* interact to control the type of marking used. If *type* is 0, simple marking is performed using the mark character(s) found in the *mark* argument. If *type* is greater than 0, automatic numbering or alphabetizing is done, and *mark* is then interpreted as the first item in the sequence to be used for numbering or alphabetizing, i.e., it is chosen from the set (1, A, a, I, i) as in {5.3.3.1}. That is:

<i>Type</i>	<i>Mark</i>	<i>Result</i>
0	omitted	hanging indent
0	<i>string</i>	<i>string</i> is the mark
>0	omitted	arabic numbering
>0	one of: 1, A, a, I, i	automatic numbering or alphabetic sequencing

Each non-zero value of *type* from 1 to 6 selects a different way of displaying the items. The following table shows the output appearance for each value of *type*:

<i>Type</i>	<i>Appearance</i>
1	x.
2	x)
3	(x)
4	[x]
5	<x>
6	{x}

where *x* is the generated number or letter.

■ The mark must not contain ordinary (paddable) spaces, because alignment of items will be lost if the right margin is justified {3.3}.

LI-space gives the number of blank lines (halves of a vertical space) that should be output by each *.LI* macro in the list. If omitted, *LI-space* defaults to 1; the value 0 can be used to obtain compact lists. If *LI-space* is greater than 0, the *.LI* macro issues a *.ne* request for two lines just before printing the mark.

LB-space, the number of blank lines (½ a vertical space) to be output by *.LB* itself, defaults to 0 if omitted.

There are three reasonable combinations of *LI-space* and *LB-space*. The normal case is to set *LI-space* to 1 and *LB-space* to 0, yielding one blank line *before* each item in the list; such a list is usually terminated with a *“.LE 1”* to end the list with a blank line. In the second case, for a more compact list, set *LI-space* to 0 and *LB-space* to 1, and, again, use *“.LE 1”* at the end of the list. The result is a list with one blank line before and after it. If you set both *LI-space* and *LB-space* to 0, and use *“.LE”* to end the list, a list without *any* blank lines will result.

Appendix A shows the definitions of the list-initialization macros {5.3.3} in terms of the *.LB* macro. Appendix B illustrates how the user can build upon those macros to obtain other kinds of lists.

6. MEMORANDUM AND RELEASED PAPER STYLES

One use of PWB/MM is for the preparation of memoranda and released papers, which have special requirements for the first page and for the cover sheet. The information needed for the memorandum or released paper (title, author, date, case numbers, etc.) is entered in the same way for *both* styles; an argument to one macro indicates which style is being used. The following sections describe the macros used to provide this data. The required order is shown in {6.9}.

8. The *mark-indent* argument is typically 0.

If neither the memorandum nor released-paper style is desired, the macros described below should be omitted from the input text. If these macros are omitted, the first page will simply have the page header {9} followed by the body of the document.

6.1 Title

.TL [charging-case] [filing-case]
one or more lines of title text

The arguments to the .TL macro are the charging case number(s) and filing case number(s).⁹ The title of the memorandum or paper follows the .TL macro and is processed in fill mode {3.1}. Multiple charging case numbers are entered as "sub-arguments" by separating each from the previous with a comma and a space, and enclosing the *entire* argument within double quotes. Multiple filing case numbers are entered similarly. For example:

.TL "12345-678, 31415-999" 31415-000
On the construction of a table
of all even prime numbers

The .br request may be used to break the title into several lines.

On output, the title appears after the word "subject" in the memorandum style. In the released-paper style, the title is centered and underlined (bold).

6.2 Author(s)

.AU name [initials] [loc] [dept] [ext] [room] [arg] [arg] [arg]

The .AU macro receives as arguments information that describes an author. If any argument contains blanks, it must be enclosed within double quotes. The first six arguments must appear in the order given (a separate .AU macro is required for each author). For example:

.AU "J. J. Jones" JJJ PY 9876 5301 9J-123

In the "from" portion in the memorandum style, the author's name is followed by location and department number on one line and by room number and extension number on the next. The "x" for the extension is added automatically. The printing of the location, department number, extension number, and room number may be suppressed on the first page of a memorandum by setting the register *Au* to 0; the default value for *Au* is 1. Arguments 7 through 9, if present, will follow this "normal" author information, each on a separate line. Certain departments have their own numbering schemes for memoranda, engineer's notes, etc. These departmental numbers usually appear after the author's name. This can be done by providing more than six arguments to the .AU macro, e.g.:

.AU "S. P. Lename" SPL IH 9988 7766 5H-444 9988.123-1MF

The name, initials, location, and department are also used in the Signature Block {6.11.1}. The author information in the "from" portion, as well as the names and initials in the Signature Block will appear in the same order as the .AU macros.

The names of the authors in the released-paper style are centered below the title. After the name of the last author, "Bell Laboratories" and the location are centered. For the case of authors from different locations, see {6.8}.

6.3 TM Number(s)

.TM [number] ...

If the memorandum is a Technical Memorandum, the TM numbers are supplied via the .TM macro. Up to nine numbers may be specified. Example:

.TM 75-9999-13 75-9988-1

This macro call is ignored in the released-paper and external-letter styles {6.6}.

9. The "charging case" is the case number to which time was charged for the development of the project described in the memorandum. The "filing case" is a number under which the memorandum is to be filed.

6.4 Abstract

.AS [arg] [indent]
text of the abstract
.AE

In both the memorandum and released-paper styles, the text of the abstract follows the author information and is preceded by the centered and underlined (*italic*) word "ABSTRACT."

The .AS (abstract start) and .AE (abstract end) macros bracket the (optional) abstract. The first argument to .AS controls the printing of the abstract. If it is 0 or null, the abstract is printed on the first page of the document, immediately following the author information, and is also saved for the cover sheet. If the first argument is 1, the abstract is saved and printed only on the cover sheet. The margins of the abstract are indented on the left and right by five spaces. The amount of indentation can be changed by specifying the desired indentation as the second argument.¹⁰

Note that headings {4.2, 4.3}, displays {7}, and footnotes {8} are *not* (as yet) permitted within an abstract.

6.5 Other Keywords

.OK [keyword] ...

Topical keywords should be specified on a Technical Memorandum cover sheet. Up to nine such keywords or keyword phrases may be specified as arguments to the .OK macro; if any keyword contains spaces, it must be enclosed within double quotes.

6.6 Memorandum Types

.MT [type] [1]

The .MT macro controls the format of the top part of the first page of a memorandum or of a released paper, as well as the format of the cover sheets. Legal codes for *type* and the corresponding values are:

<i>Code</i>	<i>Value</i>
.MT ""	no memorandum type is printed
.MT 0	no memorandum type is printed
.MT	MEMORANDUM FOR FILE
.MT 1	MEMORANDUM FOR FILE
.MT 2	PROGRAMMER'S NOTES
.MT 3	ENGINEER'S NOTES
.MT 4	Released-Paper style
.MT 5	External-Letter style
.MT "string"	string

If *type* indicates a memorandum style, then *value* will be printed after the last line of author information or after the last line of the abstract, if one appears on the first page. If *type* is longer than one character, then it, itself, will be printed. For example:

.MT "Technical Note #5"

A simple letter is produced by calling .MT with a null (but *not* omitted!) or zero argument.

The second argument to .MT is used only if the first argument is 4 (i.e., for the released-paper style) as explained in {6.8}.

In the external-letter style (.MT 5), only the date is printed in the upper right corner of the first page. It is expected that preprinted stationery will be used, providing the author's company logotype and address.

¹⁰ Values that specify indentation must be *unscaled* and are treated as "character positions," i.e., as the number of *ens*.

6.7 Date and Format Changes

6.7.1 Changing the Date. By default, the current date appears in the "date" part of a memorandum. **This can be overridden** by using:

.ND new-date

The .ND macro alters the value of the string *DT*, which is initially set to the current date.

6.7.2 Alternate First-Page Format. One can specify that the words "subject," "date," and "from" (in the memorandum style) be omitted and that an alternate company name be used:

.AF [company-name]

If an argument is given, it replaces "Bell Laboratories", without affecting the other headings. If the argument is *null*, "Bell Laboratories" is suppressed; in this case, extra blank lines are inserted to allow room for stamping the document with a Bell System logo or a Bell Laboratories stamp. .AF with *no* argument suppresses "Bell Laboratories" and the "Subject/Date/From" headings, thus allowing output on preprinted stationery.

The only .AF option appropriate for *troff* is to specify an argument to replace "Bell Laboratories" with another name.

6.8 Released-Paper Style

The released-paper style is obtained by specifying:

.MT 4 [1]

This results in a centered, underlined (bold) title followed by centered names of authors. The location of the last author is used as the location following "Bell Laboratories" (unless .AF {6.7.2} specifies a different company). If the optional second argument to .MT is given, then the name of each author is followed by the respective company name and location. The abstract, if present, follows the author information.

Information necessary for the memorandum style but not for the released-paper style is ignored.

If the released-paper style is utilized, most BTL location codes¹¹ are defined as strings that are the addresses of the corresponding BTL locations. These codes are needed only until the .MT macro is invoked. Thus, *following* the .MT macro, the user may re-use these string names. In addition, the macros described in {6.11} and their associated lines of input are ignored when the released-paper style is specified.

Authors from non-BTL locations may include their affiliations in the released-paper style by specifying the appropriate .AF *before* each .AU. For example:

```
.TL
A Learned Treatise
.AF "Getem Inc."
.AU "F. Swatter"
.AF "Bell Laboratories"
.AU "Sam P. Lename" "" CB
.MT 4 1
```

6.9 Order of Invocation of "Beginning" Macros

The macros described in {6.1-6.7}, *if present*, must be given in the following order:

11. The complete list is: AK, CP, CH, CB, DR, HO, IN, IH, MV, MH, PY, RR, RD, WV, and WH.


```
.ND new-date
.TL [charging-case] [filing-case]
one or more lines of text
.AF [company-name]
.AU name [initials] [loc] [dept] [ext] [room] [arg] [arg] [arg]
.TM [number] ...
.AS [arg] [indent]
one or more lines of text
.AE
.OK [keyword] ...
.MT [type] [1]
```

The only *required* macros for a memorandum or a released paper are .TL, .AU, and .MT; all the others (and their associated input lines) may be omitted if the features they provide are not needed. Once .MT has been invoked, *none* of the above macros can be re-invoked because they are removed from the table of defined macros to save space.

6.10 Example

The input text for this manual begins as follows:

```
.TL
P\s-3WB/MM\s0\emProgrammer's Workbench Memorandum Macros
.AU "D. W. Smith" DWS PY ...
.AU "J. R. Mashey" JRM MH ...
.MT 4 1
```

6.11 Macros for the End of a Memorandum

At the end of a memorandum (but not of a released paper), the signatures of the authors and a list of notations¹² can be requested. The following macros and their input are ignored if the released-paper style is selected.

6.11.1 Signature Block.

```
.SG [arg] [1]
```

.SG prints the author name(s) after the last line of text, aligned with the "Date/From" block. Three blank lines are left above each name for the actual signature. If no argument is given, the line of reference data¹³ will *not* appear following the last line.

A non-null first argument is treated as the typist's initials, and is appended to the reference data. Supply a null argument to print reference data with neither the typist's initials nor the preceding hyphen.

If there are several authors and if the second argument is given, then the reference data is placed on the same line as the name of the first author, rather than on the line that has the name of the last author.

The reference data contains only the location and department number of the first author. Thus, if there are authors from different departments and/or from different locations, the reference data should be supplied manually after the invocation (without arguments) of the .SG macro. For example:

```
.SG
.rs
.sp -1v
PY/MH-9876/4321-JJJ/SPL-cen
```

12. See [2], pp 1.12-16

13. The following information is known as reference data: location code, department number, author's initials, and typist's initials, all separated by hyphens. See [2], page 1.11

6.11.2 "Copy to" and Other Notations.

.NS [arg]
 zero or more lines of the notation
 .NE

After the signature and reference data, many types of notations may follow, such as a list of attachments or "copy to" lists. The various notations are obtained through the .NS macro, which provides for the proper spacing and for breaking the notations across pages, if necessary.

The codes for *arg* and the corresponding notations are:

<i>Code</i>	<i>Notations</i>
.NS ""	Copy to
.NS 0	Copy to
.NS	Copy to
.NS 1	Copy (with att.) to
.NS 2	Copy (without att.) to
.NS 3	Att.
.NS 4	Atts.
.NS 5	Enc.
.NS 6	Encs.
.NS 7	Under Separate Cover
.NS 8	Letter to
.NS 9	Memorandum to
.NS " <i>string</i> "	Copy (<i>string</i>) to

If *arg* consists of more than one character, it is placed within parentheses between the words "Copy" and "to." For example:

.NS "with att. 1 only"

will generate "Copy (with att. 1 only) to" as the notation. More than one notation may be specified before the .NE occurs, because a .NS macro terminates the preceding notation, if any. For example:

```
.NS 4
Attachment 1-List of register names
Attachment 2-List of string and macro names
.NS 1
J. J. Jones
.NS 2
S. P. Lename
G. H. Hurtz
.NE
```

would be formatted as:

```
Atts.
Attachment 1-List of register names
Attachment 2-List of string and macro names

Copy (with att.) to
J. J. Jones

Copy (without att.) to
S. P. Lename
G. H. Hurtz
```

6.12 Forcing a One-Page Letter

At times, one would like just a bit more space on the page, forcing the signature or items within notations onto the bottom of the page, so that the letter or memo is just one page in length. This can be accomplished by increasing the page length through the -rLn option, e.g. -rL90. This has the effect of

making the formatter believe that the page is 90 lines long and therefore giving it more room than usual to place the signature or the notations. This will *only* work for a *single-page* letter or memo.

7. DISPLAYS

Displays are blocks of text that are to be kept together—not split across pages. PWB/MM provides two styles of displays:¹⁴ a *static* (.DS) style and a *floating* (.DF) style. In the *static* style, the display appears in the same relative position in the output text as it does in the input text; this may result in extra white space at the bottom of the page if the display is too big to fit there. In the *floating* style, the display “floats” through the input text to the top of the next page if there is not enough room for it on the current page; thus the input text that *follows* a floating display may *precede* it in the output text. A queue of floating displays is maintained so that their relative order is not disturbed.

By default, a display is processed in no-fill mode and is *not* indented from the existing margin. The user can specify indentation or centering, as well as fill-mode processing.

Displays and footnotes {8} may *never* be nested, in any combination whatsoever. Although lists {5} and paragraphs {4.1} are permitted, no headings (.H or .HU) can occur within displays or footnotes.

7.1 Static Displays

```
.DS [format] [fill]
one or more lines of text
.DE
```

A static display is started by the .DS macro and terminated by the .DE macro. With no arguments, .DS will accept the lines of text exactly as they are typed (no-fill mode) and will *not* indent them from the prevailing indentation. The *format* argument to .DS is an integer with the following meanings:

Code	Meaning
""	no indent
0	no indent
1	indent by standard amount
2	center each line

The *fill* argument is also an integer and can have the following meanings:

Code	Meaning
""	no-fill mode
0	no-fill mode
1	fill mode

Omitted arguments are taken to be zero.

The standard amount of indentation is taken from the register *S_i*, which is initially 5. Thus, by default, the text of an indented display aligns with the first line of indented paragraphs, whose indent is contained in the *P_i* register {4.1}. Even though their initial values are the same, these two registers are independent of one another.

By default, a blank line (½ a vertical space) is placed before and after static and floating displays. These blank lines before and after *static* displays can be inhibited by setting the register *D_s* to 0.

7.2 Floating Displays

```
.DF [format] [fill]
one or more lines of text
.DE
```

A floating display is started by the .DF macro and terminated by the .DE macro. The arguments have the same meanings as for .DS {7.1}, except that, for floating displays, indent, no indent, and centering are always calculated with respect to the initial left margin, because the prevailing indent may change

14. Displays are processed in an environment that is different from that of the body of the text (see the .ev request in [9]).

between the time when the formatter first reads the floating display and the time that the display is printed. One blank line (½ a vertical space) *always* occurs both before and after a floating display.

7.3 Tables

```
.DS
.TS
one or more lines of text to be processed by tbl(I)
.TE
.DE
```

The .TS (table start) and .TE (table end) macros make possible the use of the *tbl(I)* processor [11]. They are used *only* to delimit the text to be examined by *tbl(I)*. Thus, the display function and the *tbl(I)* delimiting function are independent of one another, in order to permit one to keep together blocks that contain any mixture of tables, equations, filled and unfilled text, and caption lines.

If a particular document does not need this flexibility, it is possible to define .TS and .TE so that they act like .DS and .DE, respectively, and are also recognized by *tbl(I)*:

```
.de TS
.DS \\$1 \\$2
..
.de TE
.DE
..
```

If floating tables are desired, substitute .DF for .DS in the above.

7.4 Equations

```
.DS
.EQ
equation(s)
.EN
.DE
```

The equation setters *eqn(I)* and *neqn(I)* [6,7] expect to use the .EQ (equation start) and .EN (equation end) macros as delimiters in the same way that *tbl(I)* uses .TS and .TE; .EQ and .EN must occur either inside .DS-.DE pairs or else be defined by the user as shown above for the .TS and .TE macros.

■ *There is an exception to this rule: if .EQ and .EN are used only to specify the delimiters for in-line equations or to specify eqn/neqn "defines," .DS and .DE must not be used; otherwise extra blank lines will appear in the output.*

7.5 Figure, Table, and Equation Captions

```
.FG [title] [override] [flag]
.TB [title] [override] [flag]
.EC [title] [override] [flag]
```

The .FG (Figure Title), .TB (Table Title), .EC (Equation Caption) macros are normally used inside .DS-.DE pairs to automatically number and title figures, tables, and equations. They use registers *Fg*, *Tb*, and *Ec*, respectively.¹⁵ As an example, the call:

```
.FG "This is an illustration"
```

yields:

Figure 1. This is an illustration

.TB replaces "Figure" by "TABLE"; .EC replaces "Figure" by "Equation". Output is centered if it can fit on a single line; otherwise, all lines but the first are indented to line up with the first character of the title. The format of the numbers may be changed using the .af request of the formatter.

¹⁵ The user may wish to reset these registers after each first-level heading [4.6].

The *override* string is used to modify the normal numbering. If *flag* is omitted or 0, *override* is used as a prefix to the number; if *flag* is 1, *override* is used as a suffix; and if *flag* is 2, *override* replaces the number. For example, to produce figures numbered within sections, supply \n(H1 for *override* on each .FG call, and reset *Fg* at the beginning of each section, as shown in {4.6}.

As a matter of style, table headings are usually placed ahead of the text of the tables, while figure and equation captions usually occur after the corresponding figures and equations.

7.6 Blocks of Filled Text

One can obtain blocks of filled text through the use of .DS or .DF. However, to have the block of filled text *centered* within the current line length, the *tbl(I)* program may be used:

```

:
:
.DS 0 1
.TS
center;
lw40
T{
:
:
T}
.TE
.DE
:
:
```

The “.DS 0 1” begins a non-indented, filled display. The *tbl(I)* parameters set up a centered table with a column width of 40 ens. The “T{ ... T)” sequence allows filled text to be input as data within a table.

8. FOOTNOTES

There are two macros that delimit the text of footnotes,¹⁶ a string used to automatically number the footnotes, and a macro that specifies the style of the footnote text.

8.1 Automatic Numbering of Footnotes

Footnotes may be automatically numbered by typing the three characters “*F” immediately after the text to be footnoted, without any intervening spaces. This will place the next sequential footnote number (in a smaller point size) a half-line above the text to be footnoted.

8.2 Delimiting Footnote Text

There are two macros that delimit the text of each footnote:

```

.FS [label]
one or more lines of footnote text
.FE
```

The .FS (footnote start) marks the beginning of the text of the footnote, and the .FE marks its end. The *label* on the .FS, if present, will be used to mark the footnote text. Otherwise, the number retrieved from the string F will be used. Note that automatically-numbered and user-labeled footnotes may be intermixed. If a footnote is labeled (.FS *label*), the text to be footnoted *must* be followed by *label*, rather than by “*F”. The text between .FS and .FE is processed in fill mode. Another .FS, a .DS, or a .DF are *not* permitted between the .FS and .FE macros. Examples:

16. Footnotes are processed in an environment that is different from that of the body of the text (see the .ev request in [9]).

1. Automatically-numbered footnote:

This is the line containing the word*F
.FS
This is the text of the footnote.
.FE
to be footnoted.

2. Labelled footnote:

This is a labeled*
.FS *
The footnote is labeled with an asterisk.
.FE
footnote.

The text of the footnote (enclosed within the .FS-.FE pair) should *immediately* follow the word to be footnoted in the input text, so that “*F” or *label* occurs at the end of a line of input and the next line is the .FS macro call. It is also good practice to append a unpaddingable space [3.3] to “*F” or *label* when they follow an end-of-sentence punctuation mark (i.e., period, question mark, exclamation point).

Appendix C illustrates the various available footnote styles as well as numbered and labeled footnotes.

8.3 Format of Footnote Text •

.FD [arg] [1]

Within the footnote text, the user can control the formatting style by specifying text hyphenation, right margin justification, and text indentation, as well as left- or right-justification of the label when text indenting is used. The .FD macro is invoked to select the appropriate style. The first argument is a number from the left column of the following table. The formatting style for each number is given by the remaining four columns. For further explanation of the first two of these columns, see the definitions of the .ad, .hy, .na, and .nh requests in [9].

0	.nh	.ad	text indent	label left justified
1	.hy	.ad	"	"
2	.nh	.na	"	"
3	.hy	.na	"	"
4	.nh	.ad	no text indent	"
5	.hy	.ad	"	"
6	.nh	.na	"	"
7	.hy	.na	"	"
8	.nh	.ad	text indent	label right justified
9	.hy	.ad	"	"
10	.nh	.na	"	"
11	.hy	.na	"	"

If the first argument to .FD is out of range, the effect is as if .FD 0 were specified. If the first argument is omitted or null, the effect is equivalent to .FD 10 in *nroff* and to .FD 0 in *troff*; these are also the respective initial defaults.

If a second argument is specified, then whenever a first-level heading is encountered, automatically-numbered footnotes begin again with 1. This is most useful with the “section-page” page numbering scheme. As an example, the input line:

.FD "" 1

maintains the default formatting style and causes footnotes to be numbered afresh after each first-level heading.

For long footnotes that continue onto the following page, it is possible that, if hyphenation is permitted, the last line of the footnote on the current page will be hyphenated. Except for this case (over which the user has control by specifying an *even* argument to .FD), hyphenation across pages is inhibited by PWB/MM.

Footnotes are separated from the body of the text by a short rule. Footnotes that continue to the next page are separated from the body of the text by a full-width rule. In *troff*, footnotes are set in type that is two points smaller than the point size used in the body of the text.

8.4 Spacing between Footnote Entries

Normally, one blank line (a three-point vertical space) separates the footnotes when more than one occurs on a page. To change this spacing, set the register *Fs* to the desired value. For example:

```
.nr Fs 2
```

will cause two blank lines (a six-point vertical space) to occur between footnotes.

9. PAGE HEADERS AND FOOTERS

Text that occurs at the top of each page is known as the *page header*. Text printed at the bottom of each page is called the *page footer*. There can be up to three lines of text associated with the header: every page, even page only, and odd page only. Thus the page header may have up to two lines of text: the line that occurs at the top of every page and the line for the even- or odd-numbered page. The same is true for the page footer.

This section first describes the default appearance of page headers and page footers, and then the ways of changing them. We use the term *header* (not qualified by *even* or *odd*) to mean the line of the page header that occurs on every page, and similarly for the term *footer*.

9.1 Default Headers and Footers

By default, each page has a centered page number as the header {9.2}. There is no default footer and no even/odd default headers or footers, except as specified in {9.9}.

With the memorandum or released-paper style, the page header on the first page is suppressed *provided* a break does not occur before .MT is called. The macros and text of {6.9} and of {9} as well as .nr and .ds requests do *not* cause a break and are permitted before the .MT macro call.

9.2 Page Header

```
.PH [arg]
```

For this and for the .EH, .OH, .PF, .EF, .OF macros, the argument is of the form:

```
"'left-part'center-part'right-part'"
```

If it is inconvenient to use the apostrophe (') as the delimiter (i.e., because it occurs within one of the parts), it may be replaced *uniformly* by *any* other character. On output, the parts are left-justified, centered, and right-justified, respectively. See {9.11} for examples.

The .PH macro specifies the header that is to appear at the top of every page. The initial value (as stated in {9.1}) is the default centered page number enclosed by hyphens. See the top of this page for an example of this default header.

If *debug mode* is set using the flag -rD1 on the command line {2.4}, additional information, printed at the top left of each page, is included in the default header. This consists of the SCCS [10] Release and Level of PWB/MM (thus identifying the current version {11.3}), followed by the current line number within the current input file.

9.3 Even-Page Header

```
.EH [arg]
```

The .EH macro supplies a line to be printed at the top of each even-numbered page, immediately *following* the header. The initial value is a blank line.

9.4 Odd-Page Header

`.OH [arg]`

This macro is the same as `.EH`, except that it applies to odd-numbered pages.

9.5 Page Footer

`.PF [arg]`

The `.PF` macro specifies the line that is to appear at the bottom of each page. Its initial value is a blank line. If the `-rCn` flag is specified on the command line {2.4}, the type of copy *follows* the footer on a separate line. In particular, if `-rC3 (DRAFT)` is specified, then, in addition, the footer is initialized to contain the date {6.7.1}, instead of being a blank line.

9.6 Even-Page Footer

`.EF [arg]`

The `.EF` macro supplies a line to be printed at the bottom of each even-numbered page, immediately *preceding* the footer. The initial value is a blank line.

9.7 Odd-Page Footer

`.OF [arg]`

This macro is the same as `.EF`, except that it applies to odd-numbered pages.

9.8 Footer on the First Page

By default, the footer is a blank line. If, in the input text, one specifies `.PF` and/or `.OF` before the end of the first page of the document, then these lines will appear at the bottom of the first page.

The header (whatever its contents) *replaces* the footer *on the first page only* if the `-rN1` flag is specified on the command line {2.4}.

9.9 Default Header and Footer with "Section-Page" Numbering

Pages can be numbered sequentially within sections {4.5}. To obtain this numbering style, specify `-rN3` on the command line. In this case, the default *footer* is a centered "section-page" number, e.g. 3-5, and the default page header is blank.

9.10 Use of Strings and Registers in Header and Footer Macros •

String and register names may be placed in the arguments to the header and footer macros. If the value of the string or register is to be computed *when the respective header or footer is printed*, the invocation must be escaped by four (4) backslashes. This is because the string or register invocation will be processed three times:

- as the argument to the header or footer macro;
- in a formatting request within the header or footer macro;
- in a `.tl` request during header or footer processing.

For example, the page number register *P* must be escaped with four backslashes in order to specify a header in which the page number is to be printed at the right margin, e.g.:

```
.PH ""Page \\\nP"
```

creates a right-justified header containing the word "Page" followed by the page number. Similarly, to specify a footer with the "section-page" style, one specifies (see {4.2.2.5} for meaning of *HI*):

```
.PF ""- \\\n(HI-\\nP -"
```

As another example, suppose that the user arranges for the string *a/* to contain the current section heading which is to be printed at the bottom of each page. The `.PF` macro call would then be:

```
.PF ""\\*a/""
```

If only one or two backslashes were used, the footer would print a constant value for *a/*, namely, its value when the `.PF` appeared in the input text.

9.11 Header and Footer Example •

The following sequence specifies blank lines for the header and footer lines, page numbers on the outside edge of each page (i.e., top left margin of even pages and top right margin of odd pages), and "Dept. 9876" on the top inside margin of each page:

```
.PH ""  
.PF ""  
.EH ""\\n\p`Dept. 9876`"  
.OH ""Dept. 9876`\\n\p`"
```

9.12 Generalized Top-of-Page Processing •

☛ *This section is intended only for users accustomed to writing formatter macros.*

During header processing, PWB/MM invokes two user-definable macros. One, the .TP macro, is invoked in the environment (see .ev request in [9]) of the header; the other, .PX, is a user-exit macro that is invoked (without arguments) when the normal environment has been restored, and with "no-space" mode already in effect.

The effective initial definition of .TP (after the first page of a document) is:

```
.de TP  
.sp  
.tl \*( )t  
.if e `tl \*( )e  
.if o `tl \*( )o  
.sp  
..
```

The string }t contains the header, the string }e contains the even-page header, and the string }o contains the odd-page header, as defined by the .PH, .EH, and .OH macros, respectively. To obtain more specialized page titles, the user may redefine the .TP macro to cause any desired header processing [11.5]. Note that formatting done within the .TP macro is processed in an environment different from that of the body.

For example, to obtain a page header that includes three centered lines of data, say, a document's number, issue date, and revision date, one could define .TP as follows:

```
.de TP  
.sp  
.ce 3  
777-888-999  
Iss. 2, AUG 1977  
Rev. 7, SEP 1977  
.sp  
..
```

The .PX macro may be used to provide text that is to appear at the top of each page after the normal header and that may have tab stops to align it with columns of text in the body of the document.

9.13 Generalized Bottom-of-Page Processing

The facility to permit user-defined processing for the bottom of each page is *not* currently available.

10. TABLE OF CONTENTS AND COVER SHEET

The table of contents and the cover sheet for a technical memorandum are produced by invoking the .TC and .CS macros, respectively. The appropriate -rBn option [2.4] must *also* be specified on the command line. These macros should normally appear only once at the *end* of the document, after the Signature Block [6.11.1] and Notations [6.11.2] macros. They may occur in either order.

The table of contents is produced at the end of the document because the entire document must be processed before the table of contents can be generated. Similarly, the cover sheet is often not needed, and is therefore produced at the end.

10.1 Table of Contents

`.TC [slevel] [spacing] [tlevel] [tab] [head1] [head2] [head3] [head4] [head5]`

The `.TC` macro generates a table of contents containing the headings that were saved for the table of contents as determined by the value of the `Cl` register {4.4}. Note that `-rB1` or `-rB3` {2.4} must also be specified to the formatter on the command line. The arguments to `.TC` control the spacing before each entry, the placement of the associated page number, and additional text on the first page of the table of contents before the word "CONTENTS."

Spacing before each entry is controlled by the first two arguments; headings whose level is less than or equal to `slevel` will have `spacing` blank lines (halves of a vertical space) before them. Both `slevel` and `spacing` default to 1. This means that first-level headings are preceded by one blank line (½ a vertical space). Note that `slevel` does *not* control what levels of heading have been saved; the saving of headings is the function of the `Cl` register {4.4}.

The third and fourth arguments control the placement of the page number for each heading. The page numbers can be justified at the right margin with either blanks or dots ("leaders") separating the heading text from the page number, or the page numbers can follow the heading text. For headings whose level is less than or equal to `tlevel` (default 2), the page numbers are justified at the right margin. In this case, the value of `tab` determines the character used to separate the heading text from the page number. If `tab` is 0 (the default value), dots (i.e., leaders) are used; if `tab` is greater than 0, spaces are used. For headings whose level is greater than `tlevel`, the page numbers are separated from the heading text by two spaces (i.e., they are "ragged right").

All additional arguments (e.g., `head1`, `head2`, etc.), if any, are horizontally centered on the page, and precede the actual table of contents itself.

If the `.TC` macro is invoked with at most four arguments, then the user-exit macro `.TX` is invoked (without arguments) before the word "CONTENTS" is printed. By defining `.TX` and invoking `.TC` with at most four arguments, the user can specify what needs to be done at the top of the (first) page of the table of contents. For example, the following input:

```
.de TX
.ce 2
Special Application
Message Transmission
.sp 2
.in +10n
Approved: \l'3i'
.in
.sp
..
.TC
```

yields:

Special Application
Message Transmission

Approved: _____

CONTENTS

⋮

10.2 Cover Sheet

.CS [pages] [other] [total] [figs] [tbls] [refs]

The **.CS** macro generates a cover sheet in either the **TM** or released-paper style.¹⁷ All of the other information for the cover sheet is obtained from the data given before the **.MT** macro call {6.9}. If the released-paper style is used, all arguments to **.CS** are ignored. If a memorandum style is used, the **.CS** macro generates the "Cover Sheet for Technical Memorandum." The arguments provide the data that appears in the lower left corner of the **TM** cover sheet [2]: the number of pages of text, the number of other pages, the total number of pages, the number of figures, the number of tables, and the number of references.

11. MISCELLANEOUS FEATURES

11.1 Bold, Italic, and Roman

.B [bold-arg] [previous-font-arg]
.I [italic-arg] [previous-font-arg]
.R

When called without arguments, **.B** (or **.I**) changes the font to bold (or italic) in *troff*, and initiates underlining in *nroff*.¹⁸ This condition continues until the occurrence of a **.R**, when the regular roman font is restored. Thus,

```
.I
here is some text.
.R
```

yields:

here is some text.

If **.B** or **.I** is called with one argument, that argument is printed in the appropriate font (underlined in *nroff*). Then the *previous* font is restored (underlining is turned off in *nroff*). If two arguments are given to a **.B** or **.I**, the second argument is then concatenated to the first with no intervening space, but is printed in the previous font (not underlined in *nroff*). For example:

```
.I italic
text
.I right -justified
```

produces:

italic text right-justified

One can use both bold and italic fonts if one intends to use *troff*, but the *nroff* version of the output does not distinguish between bold and italic. It is probably a good idea to use **.I** only, unless bold is truly required. Note that font changes in headings are handled separately {4.2.2.4.1}.

Anyone using a terminal that cannot underline might wish to insert:

```
.rm ul
.rm cu
```

at the beginning of the document to eliminate *all* underlining.

11.2 Justification of Right Margin

.SA [arg]

.SA is used to set right-margin justification for the main body of text. Two justification flags are used: *current* and *default*. **.SA 0** sets both flags to no justification, i.e., it acts like the **.na** request. **.SA 1** is the inverse: it sets both flags to cause justification, just like the **.ad** request. However, calling **.SA**

17. But only if **-rB2** or **-rB3** has been specified on the command line.

18. For ease of explanation, in this section {11.1} *nroff* behavior is described first, the convention of {1.2} notwithstanding.

without an argument causes the *current* flag to be copied from the *default* flag, thus performing either a `.na` or `.ad`, depending on what the *default* is. Initially, both flags are set for no justification in *nroff* and for justification in *troff*.

In general, the request `.na` can be used to ensure that justification is turned off, but `.SA` should be used to restore justification, rather than the `.ad` request. In this way, justification or lack thereof for the remainder of the text is specified by inserting `.SA 0` or `.SA 1` *once* at the beginning of the document.

11.3 SCCS Release Identification

The string *RE* contains the SCCS [10] Release and Level of the current version of PWB/MM. For example, typing:

```
This is version \*(RE of the macros.
```

produces:

```
This is version 12.2 of the macros.
```

It is necessary to have this string printed at least once in output brought to the authors when requesting bug-fixing or enhancements. The easiest way to do this is to specify `-rD1 {2.4}` on the command line, which causes the string *RE* to be output as part of the page header {9.2}.

11.4 Two-Column Output

PWB/MM can print two columns on a page:

```
.2C
text and formatting requests (except another .2C)
.IC
```

The `.2C` macro begins two-column processing which continues until a `.IC` macro is encountered. In two-column processing, each physical page is thought of as containing two columnar “pages” of equal (but smaller) “page” width. Page headers and footers are *not* affected by two-column processing.

11.5 Column Headings for Two-Column Output •

■ *This section is intended only for users accustomed to writing formatter macros.*

In two-column output, it is sometimes necessary to have headers over each column, as well as headers over the entire page {9}. This is accomplished by redefining the `.TP` macro {9.12} to provide header lines both for the entire page and for each of the columns. For example:

```
.de TP
.sp 2
.tl 'Page \\nP'OVERALL''
.tl ''TITLE''
.sp
.nf
.ta 16C 31R 34 50C 65R
left→center→right→left→center→right      (where → stands for the tab character)
→first column→→second column
.fi
.sp 2
..
```

The above example will produce two lines of page header text plus two lines of headers over each column. The tab stops are for a 65-en overall line length.

11.6 Vertical Spacing

```
.SP [lines]
```

There exist several ways of obtaining vertical spacing, all with different effects.

The `.sp` request spaces the number of lines specified, *unless* “no space” (`.ns`) mode is on, in which case the request is ignored. This mode is typically set at the end of a page header in order to eliminate

spacing by a .sp or .bp request that just happens to occur at the top of a page. This mode can be turned off via the .rs (“restore spacing”) request.

The .SP macro is used to avoid the accumulation of vertical space by successive macro calls. Several .SP calls in a row produce *not* the sum of their arguments, but their maximum; i.e., the following produces only 3 blank lines:

```
.SP 2
.SP 3
.SP
```

Many PWB/MM macros utilize .SP for spacing. For example, “.LE 1” {5.3.2} immediately followed by “.P” {4.1} produces only a single blank line (½ a vertical space) between the end of the list and the following paragraph. An omitted argument defaults to one blank line (*one* vertical space). Unscaled fractional amounts are permitted; like .sp, .SP is also inhibited by the .ns request.

11.7 Skipping Pages

```
.SK [pages]
```

The .SK macro skips pages, but retains the usual header and footer processing. If *pages* is omitted, null, or 0, .SK skips to the top of the next page *unless* it is currently at the top of a page, in which case it does nothing. .SK *n* skips *n* pages. That is, .SK always positions the text that follows it at the top of a page, while .SK 1 always leaves one page that is blank except for the header and footer.

11.8 Setting Point Size and Vertical Spacing

In *troff*, the default point size (obtained from the register *S* {2.4}) is 10, with a vertical spacing of 12 points (i.e., 6 lines per inch). The prevailing point size and vertical spacing may be changed by invoking the .S macro:

```
.S [arg]
```

If *arg* is null, the *previous* point size is restored. If *arg* is negative, the point size is decremented by the specified amount. If *arg* is *signed* positive, the point size is incremented by the specified amount, and if *arg* is unsigned, it is used as the new point size; if *arg* is greater than 99, the *default* point size (10) is restored. Vertical spacing is always two points greater than the point size.¹⁹

12. ERRORS AND DEBUGGING

12.1 Error Terminations

When a macro discovers an error, the following actions occur:

- A break occurs.
- To avoid confusion regarding the location of the error, the formatter output buffer (which may contain some text) is printed.
- A short message is printed giving the name of the macro that found the error, the type of error, and the approximate line number (in the current input file) of the last processed input line.
- Processing terminates, unless the register *D* {2.4} has a positive value. In the latter case, processing continues even though the output is guaranteed to be deranged from that point on.

☛ *The error message is printed by writing it directly to the user's terminal. If an output filter, such as gsi(I), 450(I), or hp(I) is being used to post-process nroff output, the message may be garbled by being intermixed with text held in that filter's output buffer.*

☛ *If either tbl(I) or eqn(i)/neqn(I), or both are being used, and if the -olist option of the formatter causes the last page of the document not to be printed, a harmless “broken pipe” message results.*

All the error messages are explained in Appendix E.

¹⁹ Footnotes [8] are printed in a size two points smaller than the point size of the body, with an additional vertical spacing of three points between footnotes.

12.2 Disappearance of Output

This usually occurs because of an unclosed diversion (e.g., missing .FE or .DE). Fortunately, the macros that use diversions are careful about it, and they check to make sure that illegal nestings do not occur. If any message is issued about a missing .DE or .FE, the appropriate action is to search backwards from the termination point looking for the corresponding .DS, .DF, or .FS.

The following command:

```
grep -n "^\[EDFT\]\[EFNQS\]" files ...
```

prints all the .DS, .DF, .DE, .FS, .FE, .TS, .TE, .EQ, and .EN macros found in *files* ..., each preceded by its file name and the line number in that file. This listing can be used to check for illegal nesting and/or omission of these macros.

13. EXTENDING AND MODIFYING THE MACROS •

13.1 Naming Conventions

In this section, the following conventions are used to describe legal names:

- n: digit
- a: lower-case letter
- A: upper-case letter
- x: any letter or digit (any alphanumeric character)
- s: special character (any non-alphanumeric character)

All other characters are literals (i.e., stand for themselves).

Note that *request*, *macro*, and *string* names are kept by the formatters in a single internal table, so that there must be no duplication among such names. *Number register* names are kept in a separate table.

13.1.1 Names Used by Formatters.

- requests: aa (most common)
an (only one, currently: .c2)
- registers: aa (normal)
.x (normal)
.s (only one, currently: .\$)
% (page number)

13.1.2 Names Used by PWB/IMM.

- macros: AA (most common, accessible to user)
A (less common, accessible to user)
)x (internal, constant)
>x (internal, dynamic)
- strings: AA (most common, accessible to user)
A (less common, accessible to user)
]x (internal, usually allocated to specific functions throughout)
}x (internal, more dynamic usage)
- registers: Aa (most common, accessible to users)
An (common, accessible to user)
A (accessible, set on command line)
:x (mostly internal, rarely accessible, usually dedicated)
;x (internal, dynamic, temporaries)

13.1.3 *Names Used by EQNINEQN and TBL.* The equation preprocessors, *eqn(I)* and *neqn(I)*, use registers and string names of the form *nn*. The table preprocessor, *tbl(I)*, uses names of the form:

a- a+ a| nn #a ## #- #^ ^a T& TW

13.1.4 *User-Definable Names.* After the above, what is left for user extensions? To avoid problems, we suggest using names that consist either of a single lower-case letter, or of a lower-case letter followed by anything other than a lower-case letter. The following is a sample naming convention:

macros: aA
Aa
strings: a
a) (or a], or a), etc.)
registers a
aA

13.2 Sample Extensions

13.2.1 *Appendix Headings.* The following gives a simple way of generating and numbering appendices:

```
.nr Hu 1
.nr a 0
.de aH
.nr a +1
.nr P 0
.PH "" Appendix \\na - \\\\n"
.SK
.HU "\\$1"
..
```

After the above initialization and definition, each call of the form “.aH “title”” begins a new page (with the page header changed to “Appendix a - n”) and generates an unnumbered heading of *title*, which, if desired, can be saved for the table of contents. Those who wish Appendix titles to be centered must, in addition, set the register *Hc* to 1 {4.2.2.3}.

13.2.2 *Hanging Indent with Tabs.* The following example illustrates the use of the hanging-indent feature of variable-item lists {5.3.3.6}. First, a user-defined macro is built to accept four arguments that make up the *mark*. Each argument is to be separated from the previous one by a tab character; tab settings are defined later. Since the first argument may begin with a period or apostrophe, the “&” is used so that the formatter will not interpret such a line as a formatter request or macro.²⁰ The “\t” is translated by the formatter into a tab character. The “\c” is used to concatenate the line of *text* that follows the macro to the line of text built by the macro. The macro definition and an example of its use are as follows:

20. The two-character sequence “\&” is understood by the formatters to be a “zero-width” space, i.e., it causes no output characters to appear.

```

.de aX
.LI
\&\\$1\\t\\$2\\t\\$3\\t\\$4\\t\\c
..
:
.ta 9n 18n 27n 36n
.VL 36
.aX .nh off \- no
No hyphenation.
Automatic hyphenation is turned off.
Words containing hyphens
(e.g., mother-in-law) may still be split across lines.
.aX .hy on \- no
Hyphenate.
Automatic hyphenation is turned on.
.aX .hc\c none none no
Hyphenation indicator character is set to "c" or removed.
During text processing the indicator is suppressed
and will not appear in the output.
Prepending the indicator to a word has the effect
of preventing hyphenation of that word.
.LE

```

(□ stands for a space)

The resulting output is:

.nh	off	—	no	No hyphenation. Automatic hyphenation is turned off. Words containing hyphens (e.g., mother-in-law) may still be split across lines.
.hy	on	—	no	Hyphenate. Automatic hyphenation is turned on.
.hc c	none	none	no	Hyphenation indicator character is set to "c" or removed. During text processing the indicator is suppressed and will not appear in the output. Prepending the indicator to a word has the effect of preventing hyphenation of that word.

14. CONCLUSION

The following are the qualities that we have tried to emphasize in PWB/MM, in approximate order of importance:

- *Robustness in the face of error*—A user need not be an *nroff/troff* expert to use these macros. When the input is incorrect, either the macros attempt to make a reasonable interpretation of the error, or a message describing the error is produced. We have tried to minimize the possibility that a user would get cryptic system messages or strange output as a result of simple errors.
- *Ease of use for simple documents*—It is not necessary to write complex sequences of commands to produce simple documents, especially those that use one of the more common BTL formats. Reasonable default values are provided, where at all possible.
- *Parameterization*—There are many different preferences in the area of document styling. Many parameters are provided so that users can adapt the output to their respective needs over a wide range of styles.
- *Extension by moderately expert users*—We have made a strong effort to use mnemonic naming conventions and consistent techniques in the construction of the macros. Naming conventions are given so that a user can add new macros or redefine existing ones, if necessary.
- *Device independence*—The most common use of PWB/MM is to print documents on hard-copy typewriter terminals, using the *nroff* formatter. The macros can be used conveniently with both 10- and

12-pitch terminals. In addition, output can be scanned with an appropriate CRT terminal. The macros have been constructed to allow compatibility with *troff*, so that output can be produced both on typewriter-like terminals and on a phototypesetter.

- *Minimization of input*—The design of the macros attempts to minimize repetitive typing. For example, if a user wants to have a blank line after all first- or second-level headings, he or she need only set a specific parameter *once* at the beginning of a document, rather than add a blank line after each such heading.

Future releases of PWB/MM will provide additional features that are found to be useful. The authors welcome comments, suggestions, and criticisms of the macros and of this manual.

Acknowledgements. We are indebted to T. A. Dolotta for his continuing guidance during the development of PWB/MM. We also thank our many users who have provided much valuable feedback, both about the macros and about this manual. Finally, because PWB/MM often approaches the limits of what is possible with the text formatters, during the implementation of PWB/MM we have generated atypical requirements and encountered unusual problems; we thank J. F. Ossanna for his willingness to add new features to the formatters and to invent ways of having the formatters perform unusual but desired actions.

References

- [1] Dolotta, T. A., Haight, R. C., and Piskorik, E. M., eds. *PWB/UNIX User's Manual—Edition 1.0*. Bell Laboratories, May 1977.
- [2] Bell Laboratories, Methods and Systems Department. Office Guide. Internal Memorandum, Bell Laboratories, April 1972 (as revised).
- [3] Kernighan, B. W. *UNIX for Beginners*. Bell Laboratories, October 1974.
- [4] Kernighan, B. W. A Tutorial Introduction to the UNIX Text Editor. Bell Laboratories, October 1974.
- [5] Kernighan, B. W. A TROFF Tutorial. Bell Laboratories, August 1976.
- [6] Kernighan, B. W., and Cherry, L. L. *Typesetting Mathematics—User's Guide (Second Edition)*. Bell Laboratories, June 1976.
- [7] Scrocca, C. New Graphic Symbols for EQN and NEQN. Bell Laboratories, September 1976.
- [8] Smith, D. W., and Piskorik, E. M. Typing Documents with PWB/MM. Bell Laboratories, October 1977.
- [9] Ossanna, J. F. NROFF/TROFF User's Manual. Bell Laboratories, October 1976.
- [10] Bonanni, L. E., and Glasser, A. L. SCCS/PWB User's Manual. Bell Laboratories, November 1977.
- [11] Lesk, M. Tbl—A Program to Format Tables. Bell Laboratories, September 1977.



Appendix A: DEFINITIONS OF LIST MACROS •

☛ This appendix is intended only for users accustomed to writing formatter macros.

Here are the definitions of the list-initialization macros {5.3.3}:²¹

```
.de AL
.if!@\$1@@ .if!@\$1@1@ .if!@\$1@a@ .if!@\$1@A@ .if!@\$1@I@ .if!@\$1@i@ .)D "AL:bad arg:\$1"
.if \n(\$<3 \{.ie \w@\$2@=0 .)L \n(Lin 0 \n(Lin-\w@0\0.@u 1 "\$1"
.el .LB 0\$2 0 2 1 "\$1" \}
.if \n(\$>2 \{.ie \w@\$2@=0 .)L \n(Lin 0 \n(Lin-\w@0\0.@u 1 "\$1" 0 1
.el .LB 0\$2 0 2 1 "\$1" 0 1 \}
..
.de BL
.nr ;0 \n(Pi
.if \n(\$>0 .if \w@\$1@>0 .nr ;0 0\$1
.if \n(\$<2 .LB \n(;0 0 1 0 \}*BU
.if \n(\$>1 .LB \n(;0 0 1 0 \}*BU 0 1
.rr ;0
..
.de DL
.nr ;0 \n(Pi
.if \n(\$>0 .if \w@\$1@>0 .nr ;0 0\$1
.if \n(\$<2 .LB \n(;0 0 1 0 \{em
.if \n(\$>1 .LB \n(;0 0 1 0 \{em 0 1
.rr ;0
..
.de ML
.if !\n(\$ .)D "ML:missing arg"
.nr ;0 \w@\$1@u/3u/\n(.su+1u\ " get size in n's
.if !\n(\$-1 .LB \n(;0 0 1 0 "\$1"
.if \n(\$-1 .if !\n(\$-2 .LB 0\$2 0 1 0 "\$1"
.if \n(\$-2 .if !\w@\$2@ .LB \n(;0 0 1 0 "\$1" 0 1
.if \n(\$-2 .if \w@\$2@ .LB 0\$2 0 1 0 "\$1" 0 1
..
.de RL
.nr ;0 6
.if \n(\$>0 .if \w@\$1@>0 .nr ;0 0\$1
.if \n(\$<2 .LB \n(;0 0 2 4
.if \n(\$>1 .LB \n(;0 0 2 4 1 0 1
.rr ;0
..
.de VL
.if !\n(\$ .)D "VL:missing arg"
.if !\n(\$-2 .LB 0\$1 0\$2 0 0
.if \n(\$-2 .LB 0\$1 0\$2 0 0 \& 0 1
..
```

Any of these can be redefined to produce different behavior: e.g., to provide two spaces between the bullet of a bullet item and its text, redefine .BL as follows before invoking it:²²

```
.de BL
.LB 3 0 2 0 \}*BU
..
```

21. On this page, @ represents the BEL character, .)D is an internal PWB/MM macro that prints error messages, and .)L is similar to .LB, except that it expects its arguments to be scaled.

22. With this redefinition, .BL cannot have any arguments.

Appendix B: USER-DEFINED LIST STRUCTURES •

■ This appendix is intended only for users accustomed to writing formatter macros.

If a large document requires complex list structures, it is useful to be able to define the appearance for each list level only once, instead of having to define it at the beginning of each list. This permits consistency of style in a large document. For example, a generalized list-initialization macro might be defined in such a way that what it does depends on the list-nesting level list nesting in effect at the time the macro is called. Suppose that levels 1 through 5 of lists are to have the following appearance:

- A.
- [1]
-
- a)
- +

The following code defines a macro (.aL) that always begins a new list and determines the type of list according to the current list level. To understand it, you should know that the number register :g is used by the PWB/MM list macros to determine the current list level; it is 0 if there is no currently active list. Each call to a list-initialization macro increments :g, and each .LE call decrements it.

```
.de aL
\"      register g is used as a local temporary to save :g before it is changed below
.nr g \n(:g
.if \ng=0 .AL A \" give me an A.
.if \ng=1 .LB \n(Li 0 1 4 \" give me a [1]
.if \ng=2 .BL \" give me a bullet
.if \ng=3 .LB \n(Li 0 2 2 a \" give me an a)
.if \ng=4 .ML + \" give me a +
..
```

This macro can be used (in conjunction with .LI and .LE) instead of .AL, .RL, .BL, .LB, and .ML. For example, the following input:

```
.aL
.LI
first line.
.aL
.LI
second line.
.LE
.LI
third line.
.LE
```

will yield:

- A. first line.
- [1] second line.
- B. third line.

There is another approach to lists that is similar to the .H mechanism. The list-initialization, as well as the .LI and the .LE macros are all included in a single macro. That macro (called .bL below) requires an argument to tell it what level of item is required; it adjusts the list level by either beginning a new list or setting the list level back to a previous value, and then issues a .LI macro call to produce the item:

```
.de bL
.ie \n($ .nr g \\\$1 \" if there is an argument, that is the level
.el .nr g \n(:g \" if no argument, use current level
.if \\\ng-\n(:g>1 .)D \"**ILLEGAL SKIPPING OF LEVEL\" \" increasing level by more than 1
.if \\\ng>\n(:g \\\{.aL \\\ng-1 \" if g > :g, begin new list
.   nr g \n(:g) \" reset g to current level (.aL changes g)
.if \n(:g>\n(:g .LC \\\ng \" if :g > g, prune back to correct level
\"   if :g = g, stay within current list
.LI \"in all cases, get out an item
..
```

For .bL to work, the previous definition of the .aL macro must be changed to obtain the value of *g* from its argument, rather than from *:g*. Invoking .bL without arguments causes it to stay at the current list level. The PWB/MM .LC macro (List Clear) removes list descriptions until the level is less than or equal to that of its argument. For example, the .H macro includes the call “.LC 0”. If text is to be resumed at the end of a list, insert the call “.LC 0” to clear out the lists completely. The example below illustrates the relatively small amount of input needed by this approach. The input text:

```
The quick brown fox jumped over the lazy dog's back.
.bL 1
first line.
.bL 2
second line.
.bL 1
third line.
.bL
fourth line.
.LC 0
fifth line.
```

yields:

The quick brown fox jumped over the lazy dog's back.

- A. first line.
- [1] second line.
- B. third line.
- C. fourth line.
- fifth line.

Appendix C: SAMPLE FOOTNOTES

The following example illustrates several footnote styles and both labeled and automatically-numbered footnotes. The actual input for the immediately following text and for the footnotes at the bottom of this page is shown on the following page:

With the footnote style set to the *nroff* default, we process a footnote¹ followed by another one.***** Using the .FD macro, we changed the footnote style to hyphenate, right margin justification, indent, and left justify the label. Here is a footnote,² and another.[†] The footnote style is now set, again via the .FD macro, to no hyphenation, no right margin justification, no indentation, and with the label left-justified. Here comes the final one.³

1. This is the first footnote text example (.FD 10). This is the default style for *nroff*. The right margin is *not* justified. Hyphenation is *not* permitted. The text is indented, and the automatically generated label is *right*-justified in the text-indent space.

***** This is the second footnote text example (.FD 10). This is also the default *nroff* style but with a long footnote label provided by the user.

2. This is the third footnote example (.FD 1). The right margin is justified, the footnote text is indented, the label is *left*-justified in the text-indent space. Although not necessarily illustrated by this example, hyphenation is permitted. The quick brown fox jumped over the lazy dog's back.

† This is the fourth footnote example (.FD 1). The style is the same as the third footnote.

3. This is the fifth footnote example (.FD 6). The right margin is *not* justified, hyphenation is *not* permitted, the footnote text is *not* indented, and the label is placed at the beginning of the first line. The quick brown fox jumped over the lazy dog's back. Now is the time for all good men to come to the aid of their country.

.FD 10

With the footnote style set to the

.I nroff

default, we process a footnote.*F

.FS

This is the first footnote text example (.FD 10).

This is the default style for

.I nroff.

The right margin is

.I not

justified.

Hyphenation is

.I not

permitted.

The text is indented, and the automatically generated label is

.I right -justified

in the text-indent space.

.FE

followed by another one.*****\□

(□ stands for a space)

.FS *****

This is the second footnote text example (.FD 10).

This is also the default

.I nroff

style but with a long footnote label provided by the user.

.FE

.FD 1

Using the .FD macro, we changed the footnote style to hyphenate, right margin justification, indent, and left justify the label.

Here is a footnote.*F

.FS

This is the third footnote example (.FD 1).

The right margin is justified, the footnote text is indented, the label is

.I left -justified

in the text-indent space.

Although not necessarily illustrated by this example, hyphenation is permitted.

The quick brown fox jumped over the lazy dog's back.

.FE

and another.\(dg\□

.FS \(dg

This is the fourth footnote example (.FD 1).

The style is the same as the third footnote.

.FE

.FD 6

The footnote style is now set, again via the .FD macro, to no hyphenation, no right margin justification, no indentation, and with the label left-justified.

Here comes the final one.*F\□

.FS

This is the fifth footnote example (.FD 6).

The right margin is

.I not

justified, hyphenation is

.I not

permitted, the footnote text is

.I not

indented, and the label is placed at the beginning of the first line.

The quick brown fox jumped over the lazy dog's back.

Now is the time for all good men to come to the aid of their country.

.FE

Appendix D: SAMPLE LETTER

■ The `nroff` and `troff` outputs corresponding to the input text below are shown on the following pages.

.ND "November 1, 1977"

.TL 31415-26

Out-of-Hours Course Description

.AU "D. W. Stevenson" DWS PY 9876 4500 1G-221

.MT 0

.DS

J. M. Jones:

.DE

.P

Please use the following description for the Out-of-Hours course

"Document Preparation on the PWB/UNIX•

.FS *

UNIX is a Trademark of Bell Laboratories.

.FE

time-sharing system":

.P

The course is intended for clerks, typists, and others who intend to use the PWB/UNIX system for preparing documentation.

The course will cover such topics as:

.VL 18

.LI Environment:

utilizing a time-sharing computer system;

accessing the system;

using appropriate output terminals.

.LI Files:

how text is stored on the system;

directories;

manipulating files.

.LI "Text editing:"

how to enter text so that subsequent revisions are easier to make;

how to use the editing system to

add, delete, and move lines of text;

how to make corrections.

.LI "Text processing:"

basic concepts;

use of general-purpose formatting packages.

.LI "Other facilities:"

additional capabilities useful to the typist such as the

.I "typo, spell, diff,"

and

.I grep

commands and a desk-calculator package.

.LE

.SG jrm

.NS

S. P. Lename

H. O. Del

M. Hill

.NE

Bell Laboratories

subject: Out-of-Hours Course Description
Case: 334455

date: November 1, 1977

from: D. W. Stevenson
PY 9876
1X-123 x5432

J. M. Jones:

Please use the following description for the Out-of-Hours course
"Document Preparation on the PWB/UNIX* time-sharing system":

The course is intended for clerks, typists, and others who
intend to use the PWB/UNIX system for preparing documentation.
The course will cover such topics as:

Environment: utilizing a time-sharing computer system;
accessing the system; using appropriate output
terminals.

Files: how text is stored on the system; directories;
manipulating files.

Text editing: how to enter text so that subsequent revisions
are easier to make; how to use the editing sys-
tem to add, delete, and move lines of text; how
to make corrections.

Text processing: basic concepts; use of general-purpose format-
ting packages.

Other facilities: additional capabilities useful to the typist
such as the typo, spell, diff, and grep com-
mands and a desk-calculator package.

PY-9876-DWS-jrm

D. W. Stevenson

Copy to
S. P. Lename
H. O. Del
M. Hill

* UNIX is a Trademark of Bell Laboratories.



Bell Laboratories

subject. **Out-of-Hours Course Description**
Case: **334455**

date. **November 1, 1977**

from. **D. W. Stevenson**
PY 9876
1X-123 x5432

J. M. Jones:

Please use the following description for the Out-of-Hours course "Document Preparation on the PWB/UNIX* time-sharing system":

The course is intended for clerks, typists, and others who intend to use the PWB/UNIX system for preparing documentation. The course will cover such topics as:

- Environment: utilizing a time-sharing computer system; accessing the system; using appropriate output terminals.
- Files: how text is stored on the system; directories; manipulating files.
- Text editing: how to enter text so that subsequent revisions are easier to make; how to use the editing system to add, delete, and move lines of text; how to make corrections.
- Text processing: basic concepts; use of general-purpose formatting packages.
- Other facilities: additional capabilities useful to the typist such as the *typo*, *spell*, *diff*, and *grep* commands and a desk-calculator package.

PY-9876-DWS-jrm

D. W. Stevenson

Copy to
S. P. Lename
H. O. Del
M. Hill

* UNIX is a Trademark of Bell Laboratories.



Appendix E: ERROR MESSAGES

I. PWB/MM Error Messages

Each PWB/MM error message consists of a standard part followed by a variable part. The standard part is of the form:

ERROR:input line *n*:

The variable part consists of a descriptive message, usually beginning with a macro name. The variable parts are listed below in alphabetical order by macro name, each with a more complete explanation:²³

Check TL, AU, AS, AE, MT sequence The proper sequence of macros for the beginning of a memorandum is shown in {6.9}. Something has disturbed this order.

AL:bad arg:value The argument to the .AL macro is not one of 1, A, a, I, or i. The incorrect argument is shown as *value*.

CS:cover sheet too long The text of the cover sheet is too long to fit on one page. The abstract should be reduced or the indent of the abstract should be decreased {6.4}.

DS:too many displays More than 26 floating displays are active at once, i.e., have been accumulated but not yet output.

DS:missing FE A display starts inside a footnote. The likely cause is the omission (or misspelling) of a .FE to end a previous footnote.

DS:missing DE .DS or .DF occurs within a display, i.e., a .DE has been omitted or mistyped.

DE:no DS or DF active .DE has been encountered but there has not been a previous .DS or .DF to match it.

FE:no FS .FE has been encountered with no previous .FS to match it.

FS:illegal inside TL or AS .FS-.FE pair cannot be used inside the memorandum title or abstract.

FS:missing FE A previous .FS was not matched by a closing .FE, i.e., an attempt is being made to begin a footnote inside another one.

FS:missing DE A footnote starts inside a display, i.e., a .DS or .DF occurs without a matching .DE.

H:bad arg:value The first argument to .H must be a single digit from 1 to 7, but *value* has been supplied instead.

H:missing FE A heading macro (.H or .HU) occurs inside a footnote.

H:missing DE A heading macro (.H or .HU) occurs inside a display.

H:missing arg .H needs at least 1 argument.

HU:missing arg .HU needs 1 argument.

LB:missing arg(s) .LB requires at least 4 arguments.

LB:too many nested lists Another list was started when there were already 6 active lists.

LE:mismatched .LE has occurred without a previous .LB or other list-initialization macro {5.3.3}. Although this is not a fatal error, the message is issued because there almost certainly exists some problem in the preceding text.

23. This list is set up by ``.LB 37 0 2 0`` {5.4}.

LI:no lists active	.LI occurs without a preceding list-initialization macro. The latter has probably been omitted, or has been separated from the .LI by an intervening .H or .HU.
ML:missing arg	.ML requires at least 1 argument.
ND:missing arg	.ND requires 1 argument.
SA:bad arg:value	The argument to .SA (if any) must be either 0 or 1. The incorrect argument is shown as <i>value</i> .
SG:missing DE	.SG occurs inside a display.
SG:missing FE	.SG occurs inside a footnote.
SG:no authors	.SG occurs without any previous .AU macro(s).
VL:missing arg	.VL requires at least 1 argument.

II. Formatter Error Messages

Most messages issued by the formatter are self-explanatory. Those error messages over which the *user* has (some) control are listed below. Any other error messages should be reported to the local system-support group.

“Cannot open *filename*” is issued if one of the files in the list of files to be processed cannot be opened. If the filename is of the form `/usr/lib/tmac.name`, then the option `-mname` specifies an incorrect *name*. If the filename is of the form `/usr/lib/term/name`, then the *nroff* option `-Tname` is incorrect. If the filename is of the form `/usr/lib/font/xx`, then the font specified in a formatter `.fp` request is incorrect.

“Exception word list full” indicates that too many words have been specified in the hyphenation exception list (via `.hw` requests).

“Line overflow” means that the output line being generated was too long for the formatter’s line buffer. The excess was discarded. See the “Word overflow” message below.

“Out of temp file space” means that additional temporary space for macro definitions, diversions, etc. cannot be allocated. This message often occurs because of unclosed diversions (missing `.FE` or `.DE`), unclosed macro definitions (e.g., missing `“.”`), or a huge table of contents.

“Too many page numbers” is issued when the list of pages specified to the formatter `-o` option is too long.

“Too many string/macro names” is issued when the pool of string and macro names is full. Unneeded strings and macros can be deleted using the `.rm` request.

“Too many number registers” means that the pool of number register names is full. Unneeded registers can be deleted by using the `.rr` request.

“Word overflow” means that a word being generated exceeded the formatter’s word buffer. The excess characters were discarded. A likely cause for this and for the “Line overflow” message above are very long lines or words generated through the misuse of `\c` or of the `.cu` request, or very long equations produced by `eqn(I)/neqn(I)`.

Appendix F: SUMMARY OF MACROS, STRINGS, AND NUMBER REGISTERS

I. Macros

The following is an alphabetical list of macro names used by PWB/MM. The first line of each item gives the name of the macro, a brief description, and a reference to the section in which the macro is described. The second line gives a prototype call of the macro.

Macros marked with an asterisk are *not*, in general, invoked directly by the user. Rather, they are "user exits" called from inside header, footer, or other macros.

- 1C One-column processing {11.4}
.1C
- 2C Two-column processing {11.4}
.2C
- AE Abstract end {6.4}
.AE
- AF Alternate format of "Subject/Date/From" block {6.7.2}
.AF [company-name]
- AL Automatically-incremented list start {5.3.3.1}
.AL [type] [text-indent] [1]
- AS Abstract start {6.4}
.AS [arg] [indent]
- AU Author information {6.2}
.AU name [initials] [loc] [dept] [ext] [room] [arg] [arg] [arg]
- B Bold (underline in *nroff*) {11.1}
.B [bold-arg] [previous-font-arg]
- BL Bullet list start {5.3.3.2}
.BL [text-indent] [1]
- CS Cover sheet {10.2}
.CS [pages] [other] [total] [figs] [tbls] [refs]
- DE Display end {7.1}
.DE
- DF Display floating start {7.2}
.DF [format] [fill]
- DL Dash list start {5.3.3.3}
.DL [text-indent] [1]
- DS Display static start {7.1}
.DS [format] [fill]
- EC Equation caption {7.5}
.EC [title] [override] [flag]
- EF Even-page footer {9.6}
.EF [arg]
- EH Even-page header {9.3}
.EH [arg]
- EN End equation display {7.4}
.EN
- EQ Equation display start {7.4}
.EQ

FD	Footnote default format {8.3} .FD [arg] [1]
FE	Footnote end {8.2} .FE
FG	Figure title {7.5} .FG [title] [override] [flag]
FS	Footnote start {8.2} .FS [label]
H	Heading—numbered {4.2} .H level [heading-text]
HC	Hyphenation character {3.4} .HC [hyphenation-indicator]
HM	Heading mark style (Arabic or Roman numerals, or letters) {4.2.2.5} .HM [arg1] ... [arg7]
HU	Heading—unnumbered {4.3} .HU heading-text
HX *	Heading user exit X (before printing heading) {4.6} .HX dlevel rlevel heading-text
HZ *	Heading user exit Z (after printing heading) {4.6} .HZ dlevel rlevel heading-text
I	Italic (underline in <i>nroff</i>) {11.1} .I [italic-arg] [previous-font-arg]
LB	List begin {5.4} .LB text-indent mark-indent pad type [mark] [LI-space] [LB-space]
LC	List-status clear {Appendix B} .LC [list-level]
LE	List end {5.3.2} .LE [1]
LI	List item {5.3.1} .LI [mark] [1]
ML	Marked list start {5.3.3.4} .ML mark [text-indent] [1]
MT	Memorandum type {6.6} .MT [type] [1]
ND	New date {6.7.1} .ND new-date
NE	Notation end {6.11.2} .NE
NS	Notation start {6.11.2} .NS [arg]
OF	Odd-page footer {9.7} .OF [arg]
OH	Odd-page header {9.4} .OH [arg]
OK	Other keywords for TM cover sheet {6.5} .OK [keyword] ...

- P Paragraph {4.1}
.P [type]
- PF Page footer {9.5}
.PF [arg]
- PH Page header {9.2}
.PH [arg]
- PX * Page-header user exit {9.12}
.PX
- R Return to regular (roman) font (end underlining in *nroff*) {11.1}
.R
- RL Reference list start {5.3.3.5}
.RL [text-indent] [1]
- S Set *troff* point size and vertical spacing {11.8}
.S [arg]
- SA Set adjustment (right-margin justification) default {11.2}
.SA [arg]
- SG Signature line {6.11.1}
.SG [arg] [1]
- SK Skip pages {11.7}
.SK [pages]
- SP Space—vertically {11.6}
.SP [lines]
- TB Table title {7.5}
.TB [title] [override] [flag]
- TC Table of contents {10.1}
.TC [slevel] [spacing] [tlevel] [tab] [head1] [head2] [head3] [head4] [head5]
- TE Table end {7.3}
.TE
- TL Title of memorandum {6.1}
.TL [charging-case] [filing-case]
- TM Technical Memorandum number(s) {6.3}
.TM [number] ...
- TP * Top-of-page macro {9.12}
.TP
- TS Table start {7.3}
.TS
- TX * Table-of-contents user exit {10.1}
.TX
- VL Variable-item list start {5.3.3.6}
.VL text-indent [mark-indent] [1]

II. Strings

The following is an alphabetical list of string names used by PWB/MM, giving for each a brief description, section reference, and initial (default) value(s). See {1.4} for notes on setting and referencing strings.

- BU Bullet {3.7}
 nroff: ¶
 troff: •
- F Footnote numberer {8.1}
 nroff: \u\\n+(:p\d
 troff: \v'-.4m`s-3\\n+(:p\s0\v'.4m'
- DT Date (current date, unless overridden) {6.7.1}
 Month day, year (e.g., October 31, 1977)
- HF Heading font list, up to seven codes for heading levels 1 through 7 {4.2.2.4.1}
 3 3 2 2 2 2 2 (all underlined in *nroff*, and B B I I I I I in *troff*)
- RE SCCS Release and Level of PWB/MM {11.3}
 Release.Level (e.g., 12.2)

Note that if the released-paper style is used, then, in addition to the above strings, certain BTL location codes are defined as strings; these location strings are needed only until the .MT macro is called {6.8}.

III. Number Registers

This section provides an alphabetical list of register names, giving for each a brief description, section reference, initial (default) value, and the legal range of values (where [m:n] means values from m to n inclusive).

Any register having a single-character name can be set from the command line. An asterisk attached to a register name indicates that that register can be set *only* from the command line or *before* the PWB/MM macro definitions are read by the formatter {2.4, 2.5}. See {1.4} for notes on setting and referencing registers.


- A * Has the effect of invoking the .AF macro without an argument {2.4}
 0, [0:1]
- Au Inhibits printing of author's location, department, room, and extension in the "from" portion of a memorandum {6.2}
 1, [0:1]
- B * Defines table-of-contents and/or cover-sheet macros {2.4}
 0, [0:3]
- C * Copy type (Original, DRAFT, etc.) {2.4}
 0 (Original), [0:3]
- Cl Contents level (i.e., level of headings saved for table of contents) {4.4}
 2, [0:7]
- D * Debug flag {2.4}
 0, [0:1]
- Ds Static display pre- and post-space {7.1}
 1, [0:1]
- Ec Equation counter, used by .EC macro {7.5}
 0, [0:?], incremented by 1 for each .EC call.
- Ej Page-ejection flag for headings {4.2.2.1}
 0 (no eject), [0:7]
- Fg Figure counter, used by .FG macro {7.5}
 0, [0:?], incremented by 1 for each .FG call.
- Fs Footnote space (i.e., spacing between footnotes) {8.4}
 1, [0:?]
- H1-H7 Heading counters for levels 1-7 {4.2.2.5}
 0, [0:?], incremented by .H of corresponding level or .HU if at level given by register *Hu*.

H2-H7 are reset to 0 by any heading at a lower-numbered level.

- Hb Heading break level (after .H and .HU) {4.2.2.2}
2, [0:7]
- Hc Heading centering level for .H and .HU {4.2.2.3}
0 (no centered headings), [0:7]
- Hi Heading temporary indent (after .H and .HU) {4.2.2.2}
1 (indent as paragraph), [0:2]
- Hs Heading space level (after .H and .HU) {4.2.2.2}
2 (space only after .H 1 and .H 2), [0:7]
- Ht Heading type (for .H: single or concatenated numbers) {4.2.2.5}
0 (concatenated numbers: 1.1.1, etc.), [0:1]
- Hu Heading level for unnumbered heading (.HU) {4.3}
2 (.HU at the same level as .H 2), [0:7]
- Hy Hyphenation control for body of document {3.4}
1 (automatic hyphenation on), [0:1]
- L * Length of page {2.4}
66, [20:?] (11i, [2i:?] in *troff*)²⁴
- Li List indent {5.3.3.1}
5, [0:?]
- N * Numbering style {2.4}
0, [0:3]
- O * Offset of page {2.4}
0, [0:?] (0.5i, [0i:?] in *troff*)²⁴
- P Page number, managed by PWB/MM {2.4}
0, [0:?]
- Pi Paragraph indent {4.1}
5, [0:?]
- Pt Paragraph type {4.1}
2 (paragraphs indented except after headings, lists, and displays), [0:2]
- S * *Troff* default point size {2.4}
10, [6:36]
- Si Standard indent for displays {7.1}
5, [0:?]
- T * Type of *nroff* output device {2.4}
0, [0:2]
- Tb Table counter {7.5}
0, [0:?], incremented by 1 for each .TB call.
- U * Underlining style (*nroff*) for .H and .HU {2.4}
0 (continuous underline when possible), [0:1]
- W * Width of page (line and title length) {2.4}
65, [10:1365] (6.5i, [2i:7.54i] in *troff*)²⁴

24. For *nroff*, these values are *unscaled* numbers representing lines or character positions; for *troff*, these values must be *scaled*.







Typing Documents with PWB/MM

D. W. Smith and E. M. Piskorik

Bell Laboratories
Piscataway, New Jersey 08854

This guide shows several examples of documents prepared with PWB/MM, a set of general-purpose formatting macros used with the PWB/UNIX* text formatters *nroff* and *troff* (as well as with the *eqn/neqn* and *tbl* programs) to produce memoranda, letters, books, manuals, etc. References to manuals for these programs are given on p. 16.

In the examples, input is shown in this
Helvetica sans serif font.

The resulting output is shown (boxed) in this
Times Roman font.

Substitutable arguments are shown in this
Times Roman Italic font.

Square brackets ([...]) indicate that the enclosed substitutable argument is optional.

All output shown in the examples was done by *troff*; *nroff* output would look somewhat different.†

Contents

Paragraphs and Headings	2
Paragraph and Heading Parameters	2
Lists and List Types	4
Nested Lists	5
Italic, Bold, and Underlining	5
Displays	6
Footnotes	6
Simple Letter—Example	7
Technical Memorandum—Example	9
Memorandum-Style Macros	11
Two-Column Output	13
Equations	14
Tables	15
How to Get Output	16
References	16

* UNIX is a Trademark of Bell Laboratories.

† For example, what we call a "blank line" is a blank line in *nroff*, but is $\frac{1}{2}$ of a vertical space in *troff*, while headings that are underlined in *nroff* are either **bold** or *italic* in *troff*.

October 1977

Paragraphs and Headings

■ The output for the following is shown on p. 3.

.H 1 "PARAGRAPHS AND HEADINGS"

This section describes the types of paragraphs and the kinds of headings that are available.

.H 2 Paragraphs

Paragraphs are specified by the .P macro.

Usually, they are indented except after headings, lists, and displays.

The number register Pt is used to change the paragraph style.

.H 2 Headings

.H 3 "Numbered Headings"

There are seven levels of numbered headings.

Level 1 is the most major or highest;

level 7, the lowest.

.P

Headings are specified with the .H macro, whose first argument is the level of heading (1 through 7).

.P

The appearance of headings varies according to the level.

On output, level 1 headings are preceded by two blank lines; all others are preceded by one blank line.

Level 1 and level 2 headings produce stand-alone headings, underlined in

.I nroff

and bold in

.I troff.

Levels 3 through 7 are run-in and underlined (or italic).

.H 3 "Unnumbered Headings."

The macro .HU is a special case

of .H, in that no heading number is printed.

Each .HU heading has the level given by the register Hu, whose initial value is 2.

Usually, the value of that register is set to make unnumbered headings (if any) occur at the lowest heading level in a document.

Paragraph and Heading Parameters

There are many parameters that can change the output appearance of headings and paragraphs. Given below are some of these parameters, their *default* values, and their meanings (level 1 is the *most major* or *highest*, while level 7 is the *lowest*):

.nr Pi 5 paragraph-indent in characters (or ens).

.nr Pt 0 never indent paragraphs.

.nr Pt 1 always indent paragraphs.

.nr Pt 2 indent paragraphs *except* after headings, lists, and displays (*default*).

.ds HF 3 3 2 2 2 2 2

font specification for each

of the 7 heading levels:

1 indicates roman,

2 indicates italic,

3 indicates bold.

1 PARAGRAPHS AND HEADINGS

This section describes the types of paragraphs and the kinds of headings that are available.

1.1 Paragraphs

Paragraphs are specified by the .P macro. Usually, they are indented except after headings, lists, and displays. The number register Pt is used to change the paragraph style.

1.2 Headings

1.2.1 Numbered Headings. There are seven levels of numbered headings. Level 1 is the most major or highest, level 7, the lowest.

Headings are specified with the .H macro, whose first argument is the level of heading (1 through 7).

The appearance of headings varies according to the level. On output, level 1 headings are preceded by two blank lines, all others are preceded by one blank line. Level 1 and level 2 headings produce stand-alone headings, underlined in *nroff* and bold in *troff*. Levels 3 through 7 are run-in and underlined (or italic).

1.2.2 Unnumbered Headings. The macro .HU is a special case of .H, in that no heading number is printed. Each .HU heading has the level given by the register Hu, whose initial value is 2. Usually, the value of that register is set to make unnumbered headings (if any) occur at the lowest heading level in a document.

.HM 1 1 1 1 1 1 1

"marking" style for each heading level; the above yields an all-numeric marking style. Available styles are: 1, 0001, A, a, I, and i.

- .nr Hb 2 lowest heading level that is stand-alone (i. e., *not* run-in with the following text).
- .nr Hc 0 lowest heading level that is centered.
- .nr Hs 2 lowest heading level after which there is a blank line.
- .nr Ht 0 heading marks will be concatenated.
- .nr Hu 2 unnumbered headings (.HU) are equivalent to numbered headings at this level for spacing, font, and counting.
- .nr Cl 2 lowest heading level to be saved for the table of contents.
- .nr Ej 0 lowest heading level that forces the start of a new page.

<i>Default Heading Style</i>	
<i>to get:</i>	<i>type:</i>
n. HEADING TextH 1 "HEADING" Text ...
n.n Heading TextH 2 "Heading" Text ...
n.n.n Heading. TextH 3 "Heading." Text ...

Lists and List Types

All lists have a *list begin* macro, one or more *list items*—each consisting of a .LI macro followed by the *list item text*—and the *list end* macro .LE. That is, lists are typed like this:

```
list begin macro
.LI
list item text ...
.LI
list item text ...
:
.LE
```

where the *list begin macro* is one of the following:

- .AL [*type*] [*indent*] automatic list
(*type* is 1, A, a, I, or i; if omitted, defaults to 1)
- .BL [*indent*] bullet list
- .DL [*indent*] dash list
- .ML *mark* [*indent*] marked list
(*mark* is the desired mark)
- .RL [*indent*] reference list
- .VL *indent* variable list

indent is the number of characters of indentation (from the current indent) at which the list is to start; if it is optional and omitted, the default indentation for the given list style is used; *mark* will appear to the left of the indentation.

■ The output for the following is shown on p. 5.

```
.AL 1
.LI
Pencilpusher, I, and Hardwired, X.
A New Kind of Set Screw.
.I "Proc IEEE"
.B 75
(1976), 235-41.
.LI
Nails, H., and Irons, R.
Fasteners for Printed Circuit Boards.
.I "Proc ASME"
.B 123
(1974), 23-24.
.LE
```

- 1 Pencilpusher, I and Hardwired X A New Kind of Set Screw *Proc. IELL 75* (1976), 235-41.
- 2 Nails, H., and Irons, R. Fasteners for Printed Circuit Boards. *Proc. ASME 123* (1974), 23-24.

Nested Lists

This is ordinary text to show the margins of the page.

.AL 1

.LI

First-level item.

.AL a

.LI

Second-level item.

.LI

Another second-level item, but somewhat longer.

.LE

.LI

Return to previous list (and to previous value of indentation) at this point.

.LI

Another line.

.LE

.P

Now we're out of the lists and at the margin that existed at the beginning of this example.

This is ordinary text to show the margins of the page.

1. First-level item.

a. Second-level item.

b. Another second-level item, but somewhat longer.

2. Return to previous list (and to previous value of indentation) at this point.

3. Another line.

Now we're out of the lists and at the margin that existed at the beginning of this example.

Italic, Bold, and Underlining

In the examples on pp. 4 and 7, the macros .I, .B, and .R are used to change to, respectively, the italic, bold, and roman fonts in *troff*. In *nroff*, both .I and .B cause underlining until the occurrence of .R, which turns it off. A single argument given to either .I or .B results in that argument being underlined by *nroff*, or printed in the corresponding font by *troff*.

Displays

Displays are blocks of text that are to be kept together—not split across pages. A static display (.DS) appears in the same relative position in the output text as it does in the input text; this may result in extra white space at the bottom of a page if a static display is too big to fit there. A floating display (.DF), on the other hand, will “float” through the input text to the top of the next page if there is not enough room for it on the current page; thus, the text that *follows* a floating display in the input may *precede* it in the output. Displays can be positioned at the left margin, indented, or centered.

```
.DS [format] [fill]      .DF [format] [fill]
text ...                text ...
.DE                      .DE
```

where *format* and *fill* have the following meanings:

<i>format</i>		<i>fill</i>	
Code	Meaning	Code	Meaning
""	no indent	""	no fill
0	no indent	0	no fill
1	indent	1	fill
2	center		

Highland Avenue, Mountain Station,
South Orange, Maplewood, Millburn, Short Hills;

.DS 1

and now

for something

completely different

.DE

Summit, Chatham, Madison,

Convent Station, Morristown, New Providence,

Murray Hill, Berkeley Heights.

Highland Avenue, Mountain Station, South
Orange, Maplewood, Millburn, Short Hills.

and now

for something

completely different

Summit, Chatham, Madison, Convent Station,

Morristown, New Providence, Murray Hill,

Berkeley Heights.

Footnotes

Two styles of footnote marking are shown on p. 7. In the first, the asterisk is the mark placed on the footnote and the following .FS macro call, while in the second, a number is *automatically* generated to mark the footnote. The macros .FS and .FE are used to delimit the footnote text that is to appear at the bottom of the page.

Among the most important occupants of the workbench are the long-nosed pliers. Without this basic tool,

FS -
As first shown by Tiger & Leopard (1975).

FE
few assemblies could be completed. They may lack the popular appeal of the sledgehammer ...

FS
According to Panther & Lion (1977).

FE
appeal of the sledgehammer ...

Among the most important occupants of the workbench are the long-nosed pliers. Without this basic tool,* few assemblies could be completed. They may lack the popular appeal of the sledgehammer ...
* As first shown by Tiger & Leopard (1975).
1 According to Panther & Lion (1977)

Simple Letter—Example

The output for the following is shown on p. 8.

.P
ND May 1, 1977
PWB/MM Class
AU J. J. Jones JJJ PY 9999 5001 1Q-100
MT
DS

To All Students:
DE
P

There will be a class on the document preparation facilities of PWB/MM on November 15-18. This class lasts for 4 half-day (morning) sessions, each consisting of a lecture and practice exercises on the system.

The meeting rooms for the class are:
DS 1
ta 15n (n represents character positions)
Monday—4D-502 (— indicates a tab)
Tuesday—4D-502
Wednesday—2B-639
Thursday—2C-641.

DE
P

Please read the following before attending class:
DL

LI
"UNIX for Beginners,"
Sections I and II.

LI
A Tutorial Introduction to the UNIX Text Editor.
R

LE
(input example continued on the next page)



Bell Laboratories

subject: PWB/MM Class date: May 1, 1977

from: J. J. Jones
PY 9999
1Q-100 x5001

To All Students:
There will be a class on the document preparation facilities of PWB/MM on November 15-18. This class lasts for 4 half-day (morning) sessions, each consisting of a lecture and practice exercises on the system.

The meeting rooms for the class are:

Monday 4D-502
Tuesday 4D-502
Wednesday 2B-639
Thursday 2C-641.

Please read the following before attending class:
— UNIX for Beginners, Sections I and II.
— A Tutorial Introduction to the UNIX Text Editor.
These can be obtained from the Computing Information Library.

PY-9999-JJJ-ac J. J. Jones
Copy to
G. H. Hertz
S. P. LeName

(input example continued from the previous page)

.P
These can be obtained from the Computing Information Library.
.SG ae
.NS
G. H. Hertz
S. P. LeName
.NE

Technical Memorandum—Example

☛ The output for the following is shown on pp. 10-12.

.nr Pt 1
 .ND "June 29, 1977"
 .TL 12345 666666
 On Constructing a Table of All
 Even Prime Numbers
 .AU "S. P. LeName" SPL PY 9999 4000 1Z-123
 .AU "G. H. Hurtz" GHH PY 9999 4001 1Z-121
 .TM 76543210
 .AS
 .P
 This is an abstract for a technical memorandum.
 The abstract will appear on the cover
 sheet and on the first page
 .I unless
 the macro .AS has an argument of 1, in which case
 the abstract will be printed only on the cover sheet.
 The TM number appears on the cover sheet
 and on the first page.
 "Other Keywords" appear only on the cover sheet.
 .P
 The abstract may consist of one or more paragraphs;
 it must fit on the cover sheet.
 .AE
 .OK "Prime Numbers" Even
 .MT
 .H 1 "INTRODUCTORY MATERIAL"
 The first line of the body of the memorandum
 immediately follows the macro call for
 the heading (.H).
 Alternately, lower-level heading macros may follow it,
 as well as macros for lists, paragraphs, and so on.
 A brief example of a list follows:
 .AL A
 .LI
 This is the first item in an alphabetical
 list in the body of this memorandum.
 .LI
 This is the second item in the list.
 .AL 1
 .LI
 This is the first item in a (numbered) sub-list.
 .LI
 This is the second item in that sub-list.
 .LE
 .LE
 .P
 This is the second paragraph under the first heading.
 In addition to alphabetized and numbered lists, there
 are bullet lists, dash lists, variable lists, etc.
 .H 2 "First Second-Level Heading"
 This is the first paragraph under a
 second-level heading.
 Notice how that heading is numbered and
 where the heading and text are printed.
 .H 1 "SECOND FIRST-LEVEL HEADING"
 This is the first paragraph under the
 second first-level heading of the memorandum.
 (input example continued on the next page)



Bell Laboratories

subject	On Constructing a Table of All Even Prime Numbers Case: 12345 File: 666666	date	June 29, 1977
		from	S. P. LeName PY 9999 1Z-123 x4000 G. H. Hurtz PY 9999 1Z-121 x4001
		TM	76543210

ABSTRACT

This is an abstract for a technical memorandum. The abstract will appear on the cover sheet and on the first page unless the macro .AS has an argument of 1, in which case the abstract will be printed only on the cover sheet. The TM number appears on the cover sheet and on the first page. "Other Keywords" appear only on the cover sheet.

The abstract may consist of one or more paragraphs, it must fit on the cover sheet.

MEMORANDUM FOR FILE

1. INTRODUCTORY MATERIAL

The first line of the body of the memorandum immediately follows the macro call for the heading (.H). Alternately, lower-level heading macros may follow

(input example continued from the previous page)

.HU REFERENCES
 .RL
 .LI
 Pencilpusher, I., and Hardwired, X.
 A New Kind of Set Screw.
 .I "Proc. IEEE"
 .B 75
 (1976), 235-41.
 .LI
 Nails, H., and Irons, R.
 Fasteners for Printed Circuit Boards.
 .I "Proc. ASME"
 .B 123
 (1974), 23-24.
 .LE
 .SG rfg
 .NS 3
 .NS 2
 G. B. Brown
 C. P. Jones
 J. J. Smith
 .NE
 CS 2 1 3 0 0 2

it, as well as macros for lists, paragraphs, and so on. A brief example of a list follows.

- A. This is the first item in an alphabetical list in the body of this memorandum.
- B. This is the second item in the list.
1. This is the first item in a (numbered) sub-list.
 2. This is the second item in that sub-list.

This is the second paragraph under the first heading. In addition to alphabetized and numbered lists, there are bullet lists, dash lists, variable lists, etc.

1.1 First Second-Level Heading

This is the first paragraph under a second-level heading. Notice how that heading is numbered and where the heading and text are printed.

2 SECOND FIRST-LEVEL HEADING

This is the first paragraph under the second first-level heading of the memorandum.

REFERENCES

- [1] Pencilpusher, I., and Hardwired, X. A New Kind of Set Screw. *Proc. IEEE* 75 (1976), 235-41.
- [2] Nails, H., and Irons, R. Fasteners for Printed Circuit Boards. *Proc. ASME* 123 (1974), 23-24.

S. P. LeName

G. H. Hurtz

PY -9999-SPL/GHH-rfg

Att

Copy (without att.) to
G. B. Brown
C. P. Jones
J. J. Smith



Bell Laboratories

Cover Sheet for TM

The information contained herein ... not for publication ...

Title **On Constructing a Table of All Even Prime Numbers** Date **June 29, 1977**

TM **76543210**

Other Keywords **Prime Numbers Even**

Author(s) Location Ext Charging Case
S. P. LeName PY 12-123 4000 Filing Case 666666
G. H. Hurtz PY 12-121 4001

ABSTRACT

This is an abstract for a technical memorandum. The abstract will appear on the cover sheet and on the first page *unless* the macro .AS has an argument of 1, in which case the abstract will be printed only on the cover sheet. The TM number appears on the cover sheet and on the first page. "Other Keywords" appear only on the cover sheet.

The abstract may consist of one or more paragraphs; it must fit on the cover sheet.

Pages Text: 2 Other: 1 Total: 3

No Figures: 0 No Tables: 0 No Refs.: 2

Z-0000-X SEE REVERSE SIDE FOR DISTRIBUTION LIST

Memorandum-Style Macros

Macros for a memorandum-style document must be invoked in the order shown on pp. 9-10. Once the "memorandum type" (.MT) macro has been invoked, none of the macros that precede it can be used. The .MT macro controls the format of the "subject, date, from" portion of the first page of the memorandum. Different arguments to the .MT macro will produce different kinds of memoranda:

Code	Meaning
.MT ""	no memorandum type is printed
.MT 0	no memorandum type is printed
.MT	MEMORANDUM FOR FILE
.MT 1	MEMORANDUM FOR FILE
.MT 2	PROGRAMMER'S NOTES
.MT 3	ENGINEER'S NOTES
.MT 4	Released-Paper style
.MT 5	External Letter

The input and the resulting output for a simple letter are shown on pp. 7-8. Note that the .TM, .AS/.AE, and .OK macros are *not* used there, and that the .MT macro has a *null* argument (""). Documents of the type shown on pages 2-3 (essentially plain text) are produced by omitting, as well, the other "memorandum-style" macros: .ND, .TL, .AU, and .MT at the beginning of the document, and .SG, .NS/.NE, and .CS at the end.

Like the .MT macro, the notation macro (.NS) may also take different arguments to produce a variety of notations following the signature line:

Code	Meaning
.NS **	Copy to
.NS 0	Copy to
.NS	Copy to
.NS 1	Copy (with att.) to
.NS 2	Copy (without att.) to
.NS 3	Att.
.NS 4	Atts.
.NS 5	Enc.
.NS 6	Encs.
.NS 7	Under Separate Cover
.NS 8	Letter to
.NS 9	Memorandum to

If the .CS macro is included in the input file (see last line of p. 10) and if the -rB2 option is included on the command line (see p. 16), a cover sheet is generated (see p. 12). (The 6 arguments to .CS are the data for the bottom of the TM cover sheet: "Pages Text," "Other," etc.) Similarly, the .TC macro, together with the -rB1 or -rB3 option (see p. 16) generates a table of contents: .CS and .TC can occur only at the end of a document.

Two-Column Output

```
.DS 2
The Declaration of Independence
.DE
.DC
.P
```

When in the Course of human events, it becomes necessary for one people to dissolve the political bands which have connected them with another, and to assume among the powers of the earth, the separate and equal station to which the Laws of Nature and of Nature's God entitle them, a decent respect to the opinions of mankind requires that they should declare the causes which impel them to the separation.

We hold these truths to be self-evident, that all men are created equal. . . .

The Declaration of Independence	
When in the Course of human events, it becomes necessary for one people to dissolve the political bands which have connected them with another, and to assume among the powers of the earth, the separate and equal station to which the Laws of Nature and of Nature's	God entitle them, a decent respect to the opinions of mankind requires that they should declare the causes which impel them to the separation. We hold these truths to be self-evident, that all men are created equal. . . .

Equations

A stand-alone equation is built within a display.

```
.DS 2
.EQ
x sup 2 over a sup 2 = sqrt { pz sup 2 + qz + r }
.EN
.DE
```

$$\frac{x^2}{a^2} = \sqrt{pz^2 + qz + r}$$

```
.DS 1
.EQ
bold V bar sub nu = left [ pile { a above b above c } right ] + left [ matrix { col { A(11) above . above . } col { . above . above . } col { . above . above A(33) } } right ] times left [ pile { alpha above beta above gamma } right ]
.EN
.DE
```

$$\bar{V}_i = \begin{bmatrix} a \\ b \\ c \end{bmatrix} + \begin{bmatrix} A(11) & . & . \\ . & . & . \\ . & . & A(33) \end{bmatrix} \times \begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix}$$

In-line equations may appear in running text if a character has been defined to mark the left and right ends of the equation. Normally, \$ is used as that character and is so defined by typing the following three lines at the beginning of the document:

```
.EQ
delim $$
.EN
The quantities $a dot$, $b dotted$, $xi tilde times y vpc$ are the values that show . . .
```

The quantities \dot{a} , \ddot{b} , $\tilde{\xi} \times \tilde{\gamma}$ are the values that show . . .

This facility can be used for preparing text that contains subscripts and superscripts:

The quantity $a_{sub j} sup 3$ is . . .

The quantity a^3 is . . .

For more examples, see p. 15 and Reference 4.

Tables

The meanings of the key-letters describing the alignment of each entry are:

c center n numerical
r right-adjust a alphabetic subcolumn
l left-adjust s spanned

Global table options are *center*, *expand*, *box*, *allbox*, *doublebox*, and *tab* (*x*).

```
.DS
.TS
allbox ;
ci s s
c c c
n n n
AT&T Common Stock
Year—Price—Dividend
1973—46-55—2.87
4—40-53—3.24
5—45-52—3.40
6—51-59—95*
```

AT&T Common Stock		
Year	Price	Dividend
1973	46-55	2.87
4	40-53	3.24
5	45-52	3.40
6	51-59	.95*

* (first quarter only)

```
.TE
* (first quarter only)
```

```
.DE
(— indicates a tab)
```

```
.EQ
delim $$
.EN
.DS
.TS
box ;
cf2 cf2
| |
Name—Definition
```

```
—
.sp
Sine—$sin ( x ) = 1 over 2j ( e sup jx - e sup -jx )$
Zeta—$zeta ( s ) = \
sum from k=1 to inf k sup -s ****( Re`s > 1 )$
.TE
.DE
```

Name	Definition
Sine	$\sin(x) = \frac{1}{2j} (e^{jx} - e^{-jx})$
Zeta	$\zeta(s) = \sum_{k=1}^{\infty} k^{-s} \quad (\text{Re } s > 1)$

For more examples, see Reference 3.

How to Get Output

Documents with text only:

```
nroff: mm [options] files
or nroff [options] -mm files
troff: troff [options] -mm files
```

Text and equations:

```
nroff: mm -e [options] files
or neqn files | nroff [options] -mm -
troff: eqn files | troff [options] -mm -
```

Text and tables:

```
nroff: mm -t [options] files
or tbl files | nroff -mm [options] -
troff: tbl files | troff -mm [options] -
```

Text, tables, and equations:

```
nroff: mm -t -e [options] files
or tbl files | neqn | nroff [options] -mm -
troff: tbl files | eqn | troff [options] -mm -
```

The following options may be specified on the above PWB/UNIX shell command lines:

- ok, m-n print only page *k*, and pages *m* through *n*.
- rB1 include macros for the table of contents.
- rB2 include macros for the cover sheet.
- rB3 include macros for both.
- rC1 OFFICIAL FILE COPY in footer.
- rC2 DATE FILE COPY in footer.
- rC3 DRAFT in footer.
- rLn set page length to *n* lines.*
- rN1 page header at bottom of first page only.
- rN2 no page number on first page.
- rN3 section-page numbering.
- rOn set page offset to *n* characters.*
- rWn set line width to *n* characters.*


Terminal type and/or pitch are usually indicated by the -hp, -ti, -450, -300S, and/or -12 options of the *mm*(1) command, if it is used (see Reference 6); otherwise, they are specified by one of the *nroff* -Tname options.

References

1. PWBIMM—Programmer's Workbench Memorandum Macros by D. W. Smith and J. R. Mashey.
2. A Tutorial Introduction to the UNIX Text Editor by B. W. Kernighan.
3. Tbl—A Program to Format Tables by M. E. Lesk.
4. Typesetting Mathematics—User's Guide (Second Edition) by B. W. Kernighan and L. L. Cherry.
5. NROFF/TROFF User's Manual by J. F. Ossanna.
6. PWB/UNIX User's Manual—Edition 1.0 by T. A. Dolotta, R. C. Haight, and E. M. Piskorik, eds.

* For *nroff*, *n* must be an *unscaled* number representing lines or character positions. For *troff*, *n* must be *scaled*.





The PWB/UNIX* document entitled:


PWB/UNIX Tutorial

is not yet available.

* UNIX is a Trademark/Service Mark of the Bell System.







Tbl — A Program to Format Tables

M. E. Lesk

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

Tbl is a document formatting preprocessor for *troff* or *nroff* which makes even fairly complex tables easy to specify and enter. It is available on the PDP-11 UNIX* system and on Honeywell 6000 GCOS. Tables are made up of columns which may be independently centered, right-adjusted, left-adjusted, or aligned by decimal points. Headings may be placed over single columns or groups of columns. A table entry may contain equations, or may consist of several rows of text. Horizontal or vertical lines may be drawn as desired in the table, and any table or element may be enclosed in a box. For example:

1970 Federal Budget Transfers (in billions of dollars)			
State	Taxes collected	Money spent	Net
New York	22.91	21.35	-1.56
New Jersey	8.33	6.96	-1.37
Connecticut	4.12	3.10	-1.02
Maine	0.74	0.67	-0.07
California	22.29	22.42	+0.13
New Mexico	0.70	1.49	+0.79
Georgia	3.30	4.28	+0.98
Mississippi	1.15	2.32	+1.17
Texas	9.33	11.13	+1.80

September 4, 1977

* UNIX is a Trademark/Service Mark of the Bell System



Tbl — A Program to Format Tables

M. E. Lesk

Bell Laboratories
Murray Hill, New Jersey 07974

Introduction.

Tbl turns a simple description of a table into a *troff* or *nroff* [1] program (list of commands) that prints the table. *Tbl* may be used on the PDP-11 UNIX [2] system and on the Honeywell 6000 GCOS system. It attempts to isolate a portion of a job that it can successfully handle and leave the remainder for other programs. Thus *tbl* may be used with the equation formatting program *eqn* [3] or various layout macro packages [4,5,6], but does not duplicate their functions.

This memorandum is divided into two parts. First we give the rules for preparing *tbl* input; then some examples are shown. The description of rules is precise but technical, and the beginning user may prefer to read the examples first, as they show some common table arrangements. A section explaining how to invoke *tbl* precedes the examples. To avoid repetition, henceforth read *troff* as "*troff* or *nroff*."

The input to *tbl* is text for a document, with tables preceded by a ".TS" (table start) command and followed by a ".TE" (table end) command. *Tbl* processes the tables, generating *troff* formatting commands, and leaves the remainder of the text unchanged. The ".TS" and ".TE" lines are copied, too, so that *troff* page layout macros (such as the memo formatting macros [4]) can use these lines to delimit and place tables as they see fit. In particular, any arguments on the ".TS" or ".TE" lines are copied but otherwise ignored, and may be used by document layout macro commands.

The format of the input is as follows:

```
text
.TS
table
.TE
text
.TS
table
.TE
text
...
```

where the format of each table is as follows:

```
.TS
options ;
format .
data
.TE
```

Each table is independent, and must contain formatting information followed by the data to be entered in the table. The formatting information, which describes the individual columns and rows of the table, may be preceded by a few options that affect the entire table. A detailed description of tables is given in the next section.

Input commands.

As indicated above, a table contains, first, global options, then a format section describing the layout of the table entries, and then the data to be printed. The format and data are always required, but not the options. The various parts of the table are entered as follows:

- 1) **OPTIONS.** There may be a single line of options affecting the whole table. If present, this line must follow the `.TS` line immediately and must contain a list of option names separated by spaces, tabs, or commas, and must be terminated by a semicolon. The allowable options are:

center — center the table (default is left-adjust);
expand — make the table as wide as the current line length;
box — enclose the table in a box;
allbox — enclose each item in the table in a box;
doublebox — enclose the table in two boxes;
tab (x) — use *x* instead of tab to separate data items.

The *tbl* program tries to keep boxed tables on one page by issuing appropriate “need” (`.ne`) commands. These requests are calculated from the number of lines in the tables, and if there are spacing commands embedded in the input, these requests may be inaccurate; use normal *troff* procedures, such as keep-release macros, in that case. The user who must have a multi-page boxed table should use macros designed for this purpose, as explained below under ‘Usage.’

- 2) **FORMAT.** The format section of the table specifies the layout of the columns. Each line in this section corresponds to one line of the table (except that the last line corresponds to all following lines up to the next `.T&`, if any — see below), and each line contains a key-letter for each column of the table. It is good practice to separate the key letters for each column by spaces or tabs. Each key-letter is one of the following:

L or **l** to indicate a left-adjusted column entry;
R or **r** to indicate a right-adjusted column entry;
C or **c** to indicate a centered column entry;
N or **n** to indicate a numerical column entry, to be aligned with other numerical entries so that the units digits of numbers line up;
A or **a** to indicate an alphabetic subcolumn; all corresponding entries are aligned on the left, and positioned so that the widest is centered within the column (see example on page 12);
S or **s** to indicate a spanned heading, i.e. to indicate that the entry from the previous column continues across this column (not allowed for the first column, obviously); or
^ to indicate a vertically spanned heading, i.e. to indicate that the entry from the previous row continues down through this row. (Not allowed for the first row of the table, obviously).

When numerical alignment is specified, a location for the decimal point is sought. The rightmost dot (.) adjacent to a digit is used as a decimal point; if there is no dot adjoining a digit, the rightmost digit is used as a units digit; if no alignment is indicated, the item is centered in the column. However, the special non-printing character string `\&` may be used to override unconditionally dots and digits, or to align alphabetic data; this string lines up where a dot normally would, and then disappears from the final output. In the example below, the items shown at the left will be aligned (in a numerical column) as shown on the right:

13	13
4.2	4.2
26.4.12	26.4.12
abc	abc
abc\&	abc
43\&3.22	433.22
749.12	749.12

Note: If numerical data are used in the same column with wider L or r type table entries, the widest *number* is centered relative to the wider L or r items (L is used instead of I for readability; they have the same meaning as key-letters). Alignment within the numerical items is preserved. This is similar to the behavior of a type data, as explained above. However, alphabetic subcolumns (requested by the a key-letter) are always slightly indented relative to L items; if necessary, the column width is increased to force this. This is not true for n type entries.

Warning: the n and a items should not be used in the same column.

For readability, the key-letters describing each column should be separated by spaces. The end of the format section is indicated by a period. The layout of the key-letters in the format section resembles the layout of the actual data in the table. Thus a simple format might appear as:

```

c s s
l n n .

```

which specifies a table of three columns. The first line of the table contains a heading centered across all three columns; each remaining line contains a left-adjusted item in the first column followed by two columns of numerical data. A sample table in this format might be:

	Overall title	
Item-a	34.22	9.1
Item-b	12.65	.02
Items: c,d,e	23	5.8
Total	69.87	14.92

There are some additional features of the key-letter system:

Horizontal lines — A key-letter may be replaced by ‘_’ (underscore) to indicate a horizontal line in place of the corresponding column entry, or by ‘=’ to indicate a double horizontal line. If any data entry is provided for this column, it is ignored and a warning message is printed.

Vertical lines — A vertical bar may be placed between column key-letters. This will cause a vertical line between the corresponding columns of the table. A vertical bar to the left of the first key-letter or to the right of the last one produces a line at the edge of the table. If two vertical bars appear between key-letters, a double vertical line is drawn.

Space between columns — A number may follow the key-letter. This indicates the amount of separation between this column and the next column. The number normally specifies the separation in *ens* (one en is about the width of the letter ‘n’).* If the “expand” option is used, then these numbers are multiplied by a constant such that the table is as wide as the current line length. The default column separation number is 3. If the separation is changed the worst case (largest space requested) governs.

* More precisely, an en is a number of points (1 point = 1/72 inch) equal to half the current type size.

Vertical spanning — Normally, vertically spanned items extending over several rows of the table are centered in their vertical range. If a key-letter is followed by **t** or **T**, any corresponding vertically spanned item will begin at the top line of its range.

Font changes — A key-letter may be followed by a string containing a font name or number preceded by the letter **f** or **F**. This indicates that the corresponding column should be in a different font from the default font (usually Roman). All font names are one or two letters; a one-letter font name should be separated from whatever follows by a space or tab. The single letters **B**, **b**, **I**, and **i** are shorter synonyms for **fB** and **fI**. Font change commands given with the table entries override these specifications.

Point size changes — A key-letter may be followed by the letter **p** or **P** and a number to indicate the point size of the corresponding table entries. The number may be a signed digit, in which case it is taken as an increment or decrement from the current point size. If both a point size and a column separation value are given, one or more blanks must separate them.

Column width indication — A key-letter may be followed by the letter **w** or **W** and a width value in parentheses. This width is used as a minimum column width. If the largest element in the column is not as wide as the width value given after the **w**, the largest element is assumed to be that wide. If the largest element in the column is wider than the specified value, its width is used. The width is also used as a default line length for included text blocks. Normal *troff* units can be used to scale the width value; if none are used, the default is *ens*. If the width specification is a unitless integer the parentheses may be omitted. If the width value is changed in a column, the *last* one given controls.

Equal width columns — A key-letter may be followed by the letter **e** or **E** to indicate equal width columns. All columns whose key-letters are followed by **e** or **E** are made the same width. This permits the user to get a group of regularly spaced columns.

Note: The order of the above features is immaterial; they need not be separated by spaces, except as indicated above to avoid ambiguities involving point size and font changes. Thus a numerical column entry in italic font and 12 point type with a minimum width of 2.5 inches and separated by 6 *ens* from the next column could be specified as

```
np12w(2.5i)f1 6
```

Alternative notation — Instead of listing the format of successive lines of a table on consecutive lines of the format section, successive line formats may be given on the same line, separated by commas, so that the format for the example above might have been written:

```
c s s, l n n .
```

Default — Column descriptors missing from the end of a format line are assumed to be **L**. The longest line in the format section, however, defines the number of columns in the table; extra columns in the data are ignored silently.

- 3) **DATA.** The data for the table are typed after the format. Normally, each table line is typed as one line of data. Very long input lines can be broken: any line whose last character is **** is combined with the following line (and the **** vanishes). The data for different columns (the table entries) are separated by tabs, or by whatever character has been specified in the option *tabs* option. There are a few special cases:

Troff commands within tables — An input line beginning with a **'.** followed by anything but a number is assumed to be a command to *troff* and is passed through unchanged, retaining its position in the table. So, for example, space within a table may be produced by **".sp"** commands in the data.

Full width horizontal lines — An input *line* containing only the character `_` (underscore) or `=` (equal sign) is taken to be a single or double line, respectively, extending the full width of the *table*.

Single column horizontal lines — An input table *entry* containing only the character `_` or `=` is taken to be a single or double line extending the full width of the *column*. To obtain these characters explicitly in a column, either precede them by `\&` or follow them by a space before the usual tab or newline.

Vertically spanned items — An input table entry containing only the character string `\^` indicates that the table entry immediately above spans downward over this row. It is equivalent to a table format key-letter of `^`.

Text blocks — In order to include a block of text as a table entry, precede it by `T{` and follow it by `T}`. Thus the sequence

```
... T{
    block of
    text
T} ...
```

is the way to enter, as a single entry in the table, something that cannot conveniently be typed as a simple string between tabs. Note that the `T}` end delimiter must begin a line; additional columns of data may follow after a tab on the same line. See the example on page 10 for an illustration of included text blocks in a table. If more than twenty or thirty text blocks are used in a table, various limits in the *troff* program are likely to be exceeded, producing diagnostics such as 'too many string/macro names' or 'too many number registers.'

Text blocks are pulled out from the table, processed separately by *troff*, and replaced in the table as a solid block. If no line length is specified in the *block of text* itself, or in the table format, the default is to use $L \times C / (N + 1)$ where L is the current line length, C is the number of table columns spanned by the text, and N is the total number of columns in the table. The other parameters (point size, font, etc.) used in setting the *block of text* are those in effect at the beginning of the table (including the effect of the `^`.TS" macro) and any table format specifications of size and font, using the `p` and `f` modifiers to the column key-letters. Commands within the text block itself are also recognized, of course. However, *troff* commands within the table data but not within the text block do not affect that block.

Warnings: — Although any number of lines may be present in a table, only the first 200 lines are used in calculating the widths of the various columns. A multi-page table, of course, may be arranged as several single-page tables if this proves to be a problem. Other difficulties with formatting may arise because, in the calculation of column widths all table entries are assumed to be in the font and size being used when the `^`.TS" command was encountered, except for font and size changes indicated (a) in the table format section and (b) within the table data (as in the entry `\s+3\fidata\fp\s0`). Therefore, although arbitrary *troff* requests may be sprinkled in a table, care must be taken to avoid confusing the width calculations; use requests such as `.ps` with care.

- 4) **ADDITIONAL COMMAND LINES.** If the format of a table must be changed after many similar lines, as with sub-headings or summarizations, the `^`.T&" (table continue) command can be used to change column parameters. The outline of such a table input is:

```
.TS
options ;
format .
data
. . .
.T&
format .
data
.T&
format .
data
.TE
```

as in the examples on pages 9 and 12. Using this procedure, each table line can be close to its corresponding format line.

Warning: it is not possible to change the number of columns, the space between columns, the global options such as *box*, or the selection of columns to be made equal width.

Usage.

On UNIX, *tbl* can be run on a simple table with the command

```
tbl input-file |troff
```

but for more complicated use, where there are several input files, and they contain equations and *ms* memorandum layout commands as well as tables, the normal command would be

```
tbl file-1 file-2 . . . |eqn |troff -ms
```

and, of course, the usual options may be used on the *troff* and *eqn* commands. The usage for *nroff* is similar to that for *troff*, but only TELETYPE® Model 37 and Diablo-mechanism (DASI or GSI) terminals can print boxed tables.

Note that when *eqn* and *tbl* are used together on the same file *tbl* should be used first. If there are no equations within tables, either order works, but it is usually faster to run *tbl* first, since *eqn* normally produces a larger expansion of the input than *tbl*. However, if there are equations within tables (using the *delim* mechanism in *eqn*), *tbl* must be first or the output will be scrambled. Users must also beware of using equations in *n*-style columns; this is nearly always wrong, since *tbl* attempts to split numerical format items into two parts and this is not possible with equations.

Tbl limits tables to twenty columns; however, use of more than 16 numerical columns may fail because of limits in *troff*, producing the 'too many number registers' message. *Troff* number registers used by *tbl* must be avoided by the user within tables; these include two-digit names from 31 to 99, and names of the forms #x, x+, x|, ^x, and x-, where x is any lower case letter. The names ##, #-, and #^ are also used in certain circumstances. To conserve number register names, the *n* and *a* formats share a register; hence the restriction above that they may not be used in the same column.

For aid in writing layout macros, *tbl* defines a number register TW which is the table width; it is defined by the time that the ".TE" macro is invoked and may be used in the expansion of that macro. More importantly, to assist in laying out multi-page boxed tables the macro T# is defined to produce the bottom lines and side lines of a boxed table, and then invoked at its end. By use of this macro in the page footer a multi-page table can be boxed. In particular, the *ms* macros can be used to print a multi-page boxed table with a repeated heading by giving the argument H to the ".TS" macro. If the table start macro is written

```
.TS H
```

a line of the form

```
.TH
```

must be given in the table after any table heading (or at the start if none). Material up to the

“.TII” is placed at the top of each page of table; the remaining lines in the table are placed on several pages as required. Note that this is *not* a feature of *tbl*, but of the *ms* layout macros.

Examples.

Here are some examples illustrating features of *tbl*. The symbol ⊕ in the input represents a tab character.

Input:

```
.TS
box;
c c c
111.
Language ⊕ Authors ⊕ Runs on

Fortran ⊕ Many ⊕ Almost anything
PL/1 ⊕ IBM ⊕ 360/370
C ⊕ BTL ⊕ 11/45,H6000,370
BLISS ⊕ Carnegie-Mellon ⊕ PDP-10,11
IDS ⊕ Honeywell ⊕ H6000
Pascal ⊕ Stanford ⊕ 370
.TE
```

Output:

Language	Authors	Runs on
Fortran	Many	Almost anything
PL/1	IBM	360/370
C	BTL	11/45,H6000,370
BLISS	Carnegie-Mellon	PDP-10,11
IDS	Honeywell	H6000
Pascal	Stanford	370

Input:

```
.TS
allbox;
c s s
c c c
n n n.
AT&T Common Stock
Year ⊕ Price ⊕ Dividend
1971 ⊕ 41-54 ⊕ $2.60
2 ⊕ 41-54 ⊕ 2.70
3 ⊕ 46-55 ⊕ 2.87
4 ⊕ 40-53 ⊕ 3.24
5 ⊕ 45-52 ⊕ 3.40
6 ⊕ 51-59 ⊕ .95*
.TE
* (first quarter only)
```

Output:

AT&T Common Stock		
Year	Price	Dividend
1971	41-54	\$2.60
2	41-54	2.70
3	46-55	2.87
4	40-53	3.24
5	45-52	3.40
6	51-59	.95*

* (first quarter only)

Input:

```
.TS
box;
c s s
c |c |c
l |l |n.
Major New York Bridges
=
Bridge @Designer @Length
_
Brooklyn @J. A. Roebling @1595
Manhattan @G. Lindenthal @1470
Williamsburg @L. L. Buck @1600
_
Queensborough @Palmer & @1182
@ Hornbostel

_
@ @1380
Triborough @O. H. Ammann @_
@ @383

_
Bronx Whitestone @O. H. Ammann @2300
Throgs Neck @O. H. Ammann @1800
_
George Washington @O. H. Ammann @3500
.TE
```

Output:

Major New York Bridges		
Bridge	Designer	Length
Brooklyn	J. A. Roebling	1595
Manhattan	G. Lindenthal	1470
Williamsburg	L. L. Buck	1600
Queensborough	Palmer & Hornbostel	1182
Triborough	O. H. Ammann	1380
		383
Bronx Whitestone	O. H. Ammann	2300
Throgs Neck	O. H. Ammann	1800
George Washington	O. H. Ammann	3500

Input:

```
.TS
c c
np-2 |n |.
@Stack
@_
1 @46
@_
2 @23
@_
3 @15
@_
4 @6.5
@_
5 @2.1
@_
.TE
```

Output:

	Stack
1	46
2	23
3	15
4	6.5
5	2.1

Input:

```
.TS
box;
L L L
L L
L L | LB
L L _
L L L.
january @ february @ march
april @ may
june @ july @ Months
august @ september
october @ november @ december
.TE
```

Output:

january	february	march
april	may	Months
june	july	
august	september	
october	november	december

Input:

```
.TS
box;
cfB s s s.
Composition of Foods

.T&
c | c s s
c | c s s
c | c | c | c.
Food @ Percent by Weight
\ ^ _
\ ^ @ Protein @ Fat @ Carbo-
\ ^ \ ^ \ ^ @ hydrate

.T&
l | n | n | n.
Apples @ .4 @ .5 @ 13.0
Halibut @ 18.4 @ 5.2 @ . . .
Lima beans @ 7.5 @ .8 @ 22.0
Milk @ 3.3 @ 4.0 @ 5.0
Mushrooms @ 3.5 @ .4 @ 6.0
Rye bread @ 9.0 @ .6 @ 52.7
.TE
```

Output:

Composition of Foods			
Food	Percent by Weight		
	Protein	Fat	Carbo- hydrate
Apples	.4	.5	13.0
Halibut	18.4	5.2	...
Lima beans	7.5	.8	22.0
Milk	3.3	4.0	5.0
Mushrooms	3.5	.4	6.0
Rye bread	9.0	.6	52.7

Input:

```
.TS
allbox;
cfl s s
c cw(1i) cw(1i)
lp9 lp9 lp9.
New York Area Rocks
Era @ Formation @ Age (years)
Precambrian @ Reading Prong @ >1 billion
Paleozoic @ Manhattan Prong @ 400 million
Mesozoic @ T{
.na
Newark Basin, incl.
Stockton, Lockatong, and Brunswick
formations; also Watchungs
and Palisades.
T} @ 200 million
Cenozoic @ Coastal Plain @ T{
On Long Island 30,000 years;
Cretaceous sediments redeposited
by recent glaciation.
.ad
T}
.TE
```

Output:

New York Area Rocks		
Era	Formation	Age (years)
Precambrian	Reading Prong	>1 billion
Paleozoic	Manhattan Prong	400 million
Mesozoic	Newark Basin, incl. Stockton, Lockatong, and Brunswick for- mations; also Watchungs and Palisades.	200 million
Cenozoic	Coastal Plain	On Long Island 30,000 years; Cretaceous sedi- ments redepo- sited by recent glaciation.

Input:

```
.EQ
delim $$
.EN
. . .
.TS
doublebox;
c c
ll.
Name @ Definition
.sp
.vs +2p
Gamma @ $GAMMA (z) = int sub 0 sup inf t sup {z-1} e sup -t dt$
Sine @ $sin (x) = 1 over 2i ( e sup ix - e sup -ix )$
Error @ $ roman erf (z) = 2 over sqrt pi int sub 0 sup z e sup {-t sup 2} dt$
Bessel @ $ J sub 0 (z) = 1 over pi int sub 0 sup pi cos ( z sin theta ) d theta $
Zeta @ $ zeta (s) = sum from k=1 to inf k sup -s ^^ ( Re s > 1)$
.vs -2p
.TE
```

Output:

Name	Definition
Gamma	$\Gamma(z) = \int_0^{\infty} t^{z-1} e^{-t} dt$
Sine	$\sin(x) = \frac{1}{2i} (e^{ix} - e^{-ix})$
Error	$\operatorname{erf}(z) = \frac{2}{\sqrt{\pi}} \int_0^z e^{-t^2} dt$
Bessel	$J_0(z) = \frac{1}{\pi} \int_0^{\pi} \cos(z \sin \theta) d\theta$
Zeta	$\zeta(s) = \sum_{k=1}^{\infty} k^{-s} \quad (\operatorname{Re} s > 1)$

Input:

```

.TS
box, tab(:);
cb s s s s
cp-2 s s s s
c||c|c|c|c
c||c|c|c|c
r2||n2|n2|n2|n.
Readability of Text
Line Width and Leading for 10-Point Type
=

```

Line : Set : 1-Point : 2-Point : 4-Point
Width : Solid : Leading : Leading : Leading

```

9 Pica : \-9.3 : \-6.0 : \-5.3 : \-7.1
14 Pica : \-4.5 : \-0.6 : \-0.3 : \-1.7
19 Pica : \-5.0 : \-5.1 : 0.0 : \-2.0
31 Pica : \-3.7 : \-3.8 : \-2.4 : \-3.6
43 Pica : \-9.1 : \-9.0 : \-5.9 : \-8.8
.TE

```

Output:

Readability of Text				
Line Width and Leading for 10-Point Type				
Line Width	Set Solid	1-Point Leading	2-Point Leading	4-Point Leading
9 Pica	-9.3	-6.0	-5.3	-7.1
14 Pica	-4.5	-0.6	-0.3	-1.7
19 Pica	-5.0	-5.1	0.0	-2.0
31 Pica	-3.7	-3.8	-2.4	-3.6
43 Pica	-9.1	-9.0	-5.9	-8.8

Input:

.TS
 c s
 cip-2 s
 l n
 a n.
 Some London Transport Statistics
 (Year 1964)
 Railway route miles ⊕ 244
 Tube ⊕ 66
 Sub-surface ⊕ 22
 Surface ⊕ 156
 .sp .5
 .T&
 l r
 a r.
 Passenger traffic \- railway
 Journeys ⊕ 674 million
 Average length ⊕ 4.55 miles
 Passenger miles ⊕ 3,066 million
 .T&
 l r
 a r.
 Passenger traffic \- road
 Journeys ⊕ 2,252 million
 Average length ⊕ 2.26 miles
 Passenger miles ⊕ 5,094 million
 .T&
 l n
 a n.
 .sp .5
 Vehicles ⊕ 12,521
 Railway motor cars ⊕ 2,905
 Railway trailer cars ⊕ 1,269
 Total railway ⊕ 4,174
 Omnibuses ⊕ 8,347
 .T&
 l n
 a n.
 .sp .5
 Staff ⊕ 73,739
 Administrative, etc. ⊕ 5,582
 Civil engineering ⊕ 5,134
 Electrical eng. ⊕ 1,714
 Mech. eng. \- railway ⊕ 4,310
 Mech. eng. \- road ⊕ 9,152
 Railway operations ⊕ 8,930
 Road operations ⊕ 35,946
 Other ⊕ 2,971
 .TE

Output:

Some London Transport Statistics
 (Year 1964)

Railway route miles	244
Tube	66
Sub-surface	22
Surface	156
Passenger traffic – railway	
Journeys	674 million
Average length	4.55 miles
Passenger miles	3,066 million
Passenger traffic – road	
Journeys	2,252 million
Average length	2.26 miles
Passenger miles	5,094 million
Vehicles	12,521
Railway motor cars	2,905
Railway trailer cars	1,269
Total railway	4,174
Omnibuses	8,347
Staff	73,739
Administrative, etc.	5,582
Civil engineering	5,134
Electrical eng.	1,714
Mech. eng. – railway	4,310
Mech. eng. – road	9,152
Railway operations	8,930
Road operations	35,946
Other	2,971

Input:

.ps 8
.vs 10p
.TS
center box;
c s s
ci s s
c c c
IB l n.
New Jersey Representatives
(Democrats)
.sp .5
Name ⊕ Office address ⊕ Phone
.sp .5
James J. Florio ⊕ 23 S. White Horse Pike, Somerdale 08083 ⊕ 609-627-8222
William J. Hughes ⊕ 2920 Atlantic Ave., Atlantic City 08401 ⊕ 609-345-4844
James J. Howard ⊕ 801 Bangs Ave., Asbury Park 07712 ⊕ 201-774-1600
Frank Thompson, Jr. ⊕ 10 Rutgers Pl., Trenton 08618 ⊕ 609-599-1619
Andrew Maguire ⊕ 115 W. Passaic St., Rochelle Park 07662 ⊕ 201-843-0240
Robert A. Roe ⊕ U.S.P.O., 194 Ward St., Paterson 07510 ⊕ 201-523-5152
Henry Helstoski ⊕ 666 Paterson Ave., East Rutherford 07073 ⊕ 201-939-9090
Peter W. Rodino, Jr. ⊕ Suite 1435A, 970 Broad St., Newark 07102 ⊕ 201-645-3213
Joseph G. Minish ⊕ 308 Main St., Orange 07050 ⊕ 201-645-6363
Helen S. Meyner ⊕ 32 Bridge St., Lambertville 08530 ⊕ 609-397-1830
Dominick V. Daniels ⊕ 895 Bergen Ave., Jersey City 07306 ⊕ 201-659-7700
Edward J. Patten ⊕ Natl. Bank Bldg., Perth Amboy 08861 ⊕ 201-826-4610
.sp .5
.T&
ci s s
IB l n.
(Republicans)
.sp .5v
Millicent Fenwick ⊕ 41 N. Bridge St., Somerville 08876 ⊕ 201-722-8200
Edwin B. Forsythe ⊕ 301 Mill St., Moorestown 08057 ⊕ 609-235-6622
Matthew J. Rinaldo ⊕ 1961 Morris Ave., Union 07083 ⊕ 201-687-4235
.TE
.ps 10
.vs 12p

Output:

New Jersey Representatives (Democrats)		
Name	Office address	Phone
James J. Florio	23 S. White Horse Pike, Somerdale 08083	609-627-8222
William J. Hughes	2920 Atlantic Ave., Atlantic City 08401	609-345-4844
James J. Howard	801 Bangs Ave., Asbury Park 07712	201-774-1600
Frank Thompson, Jr.	10 Rutgers Pl., Trenton 08618	609-599-1619
Andrew Maguire	115 W. Passaic St., Rochelle Park 07662	201-843-0240
Robert A. Roe	U.S.P.O., 194 Ward St., Paterson 07510	201-523-5152
Henry Helstoski	666 Paterson Ave., East Rutherford 07073	201-939-9090
Peter W. Rodino, Jr.	Suite 1435A, 970 Broad St., Newark 07102	201-645-3213
Joseph G. Minish	308 Main St., Orange 07050	201-645-6363
Helen S. Meyner	32 Bridge St., Lambertville 08530	609-397-1830
Dominick V. Daniels	895 Bergen Ave., Jersey City 07306	201-659-7700
Edward J. Patten	Natl. Bank Bldg., Perth Amboy 08861	201-826-4610
(Republicans)		
Millicent Fenwick	41 N. Bridge St., Somerville 08876	201-722-8200
Edwin B. Forsythe	301 Mill St., Moorestown 08057	609-235-6622
Matthew J. Rinaldo	1961 Morris Ave., Union 07083	201-687-4235

This is a paragraph of normal text placed here only to indicate where the left and right margins are. In this way the reader can judge the appearance of centered tables or expanded tables, and observe how such tables are formatted.

Input:

```
.TS
expand;
c s s s
c c c c
l l n n.
Bell Labs Locations
Name ⊕ Address ⊕ Area Code ⊕ Phone
Holmdel ⊕ Holmdel, N. J. 07733 ⊕ 201 ⊕ 949-3000
Murray Hill ⊕ Murray Hill, N. J. 07974 ⊕ 201 ⊕ 582-6377
Whippany ⊕ Whippany, N. J. 07981 ⊕ 201 ⊕ 386-3000
Indian Hill ⊕ Naperville, Illinois 60540 ⊕ 312 ⊕ 690-2000
.TE
```

Output:

Bell Labs Locations			
Name	Address	Area Code	Phone
Holmdel	Holmdel, N. J. 07733	201	949-3000
Murray Hill	Murray Hill, N. J. 07974	201	582-6377
Whippany	Whippany, N. J. 07981	201	386-3000
Indian Hill	Naperville, Illinois 60540	312	690-2000

Input:

.TS
box;
cb s s s
c|c|c s
ltiw(1i)|lw(2i)|lp8|lw(1.6i)p8.
Some Interesting Places

NameⓄDescriptionⓄPractical Information

T{

American Museum of Natural History

T}ⓄT{

The collections fill 11.5 acres (Michelin) or 25 acres (MTA) of exhibition halls on four floors. There is a full-sized replica of a blue whale and the world's largest star sapphire (stolen in 1964).

T}ⓄHoursⓄ10-5, ex. Sun 11-5, Wed. to 9

\`Ⓞ\`ⓄLocationⓄT{

Central Park West & 79th St.

T}

\`Ⓞ\`ⓄAdmissionⓄDonation: \$1.00 asked

\`Ⓞ\`ⓄSubwayⓄAA to 81st St.

\`Ⓞ\`ⓄTelephoneⓄ212-873-4225

Bronx ZooⓄT{

About a mile long and .6 mile wide, this is the largest zoo in America.

A lion eats 18 pounds

of meat a day while a sea lion eats 15 pounds of fish.

T}ⓄHoursⓄT{

10-4:30 winter, to 5:00 summer

T}

\`Ⓞ\`ⓄLocationⓄT{

185th St. & Southern Blvd, the Bronx.

T}

\`Ⓞ\`ⓄAdmissionⓄ\$1.00, but Tu, We, Th free

\`Ⓞ\`ⓄSubwayⓄ2, 5 to East Tremont Ave.

\`Ⓞ\`ⓄTelephoneⓄ212-933-1759

Brooklyn MuseumⓄT{

Five floors of galleries contain American and ancient art.

There are American period rooms and architectural ornaments saved from wreckers, such as a classical figure from Pennsylvania Station.

T}ⓄHoursⓄWed-Sat, 10-5, Sun 12-5

\`Ⓞ\`ⓄLocationⓄT{

Eastern Parkway & Washington Ave., Brooklyn.

T}

\`Ⓞ\`ⓄAdmissionⓄFree

\`Ⓞ\`ⓄSubwayⓄ2,3 to Eastern Parkway.

\`Ⓞ\`ⓄTelephoneⓄ212-638-5000

T{

New-York Historical Society

T}ⓄT{

All the original paintings for Audubon's

.I

Birds of America

.R

are here, as are exhibits of American decorative arts, New York history, Hudson River school paintings, carriages, and glass paperweights.

T}ⓄHoursⓄT{

Tues-Fri & Sun, 1-5; Sat 10-5

T}

\`Ⓞ\`ⓄLocationⓄT{

Central Park West & 77th St.

T}

\`Ⓞ\`ⓄAdmissionⓄFree

\`Ⓞ\`ⓄSubwayⓄAA to 81st St.

\`Ⓞ\`ⓄTelephoneⓄ212-873-3400

.TE

Output:

Some Interesting Places			
Name	Description	Practical Information	
<i>American Museum of Natural History</i>	The collections fill 11.5 acres (Michelin) or 25 acres (MTA) of exhibition halls on four floors. There is a full-sized replica of a blue whale and the world's largest star sapphire (stolen in 1964).	Hours Location Admission Subway Telephone	10-5, ex. Sun 11-5, Wed. to 9 Central Park West & 79th St. Donation: \$1.00 asked AA to 81st St. 212-873-4225
<i>Bronx Zoo</i>	About a mile long and .6 mile wide, this is the largest zoo in America. A lion eats 18 pounds of meat a day while a sea lion eats 15 pounds of fish.	Hours Location Admission Subway Telephone	10-4:30 winter, to 5:00 summer 185th St. & Southern Blvd, the Bronx. \$1.00, but Tu, We, Th free 2, 5 to East Tremont Ave. 212-933-1759
<i>Brooklyn Museum</i>	Five floors of galleries contain American and ancient art. There are American period rooms and architectural ornaments saved from wreckers, such as a classical figure from Pennsylvania Station.	Hours Location Admission Subway Telephone	Wed-Sat, 10-5, Sun 12-5 Eastern Parkway & Washington Ave., Brooklyn. Free 2,3 to Eastern Parkway. 212-638-5000
<i>New-York Historical Society</i>	All the original paintings for Audubon's <i>Birds of America</i> are here, as are exhibits of American decorative arts, New York history, Hudson River school paintings, carriages, and glass paperweights.	Hours Location Admission Subway Telephone	Tues-Fri & Sun, 1-5; Sat 10-5 Central Park West & 77th St. Free AA to 81st St. 212-873-3400

Acknowledgments.

Many thanks are due to J. C. Blinn, who has done a large amount of testing and assisted with the design of the program. He has also written many of the more intelligible sentences in this document and helped edit all of it. All phototypesetting programs on UNIX are dependent on the work of J. F. Ossanna, whose assistance with this program in particular has been most helpful. This program is patterned on a table formatter originally written by J. F. Gimpel. The assistance of T. A. Dolotta, B. W. Kernighan, and J. N. Sturman is gratefully acknowledged.

References.

- [1] J. F. Ossanna, *NROFF/TROFF User's Manual*, Computing Science Technical Report No. 55, Bell Laboratories, 1976.
- [2] K. Thompson and D. M. Ritchie, "The UNIX Time-Sharing System," *Comm. ACM.* **17**, pp. 365-75 (1974).
- [3] B. W. Kernighan and L. L. Cherry, "A System for Typesetting Mathematics," *Comm. ACM.* **18**, pp. 151-57 (1975).
- [4] M. E. Lesk, *Typing Documents on UNIX*, Bell Laboratories internal memorandum.
- [5] M. E. Lesk and B. W. Kernighan, *Computer Typesetting of Technical Journals on UNIX*, Computing Science Technical Report No. 44, Bell Laboratories, July 1976.


- [6] J. R. Mashey and D. W. Smith, *PWBIMM — Programmer's Workbench Memorandum Macros*, Bell Laboratories memorandum.

List of Tbl Command Characters and Words

<i>Command</i>	<i>Meaning</i>	<i>Section</i>
a A	Alphabetic subcolumn	2
allbox	Draw box around all items	1
b B	Boldface item	2
box	Draw box around table	1
c C	Centered column	2
center	Center table in page	1
doublebox	Doubled box around table	1
e E	Equal width columns	2
expand	Make table full line width	1
f F	Font change	2
i I	Italic item	2
l L	Left adjusted column	2
n N	Numerical column	2
<i>nmn</i>	Column separation	2
p P	Point size change	2
r R	Right adjusted column	2
s S	Spanned item	2
t T	Vertical spanning at top	2
tab (x)	Change data separator character	1
T{ T }	Text block	3
w W	Minimum width value	2
.xx	Included <i>troff</i> command	3
	Vertical line	2
	Double vertical line	2
^	Vertical span	2
\^	Vertical span	3
=	Double horizontal line	2,3
-	Horizontal line	2,3







A TROFF Tutorial

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

troff is a text-formatting program for driving the Graphic Systems phototypesetter on the UNIX and GCOS operating systems. This device is capable of producing high quality text; this paper is an example of **troff** output.

The phototypesetter itself normally runs with four fonts, containing roman, italic and bold letters (as on this page), a full greek alphabet, and a substantial number of special characters and mathematical symbols. Characters can be printed in a range of sizes, and placed anywhere on the page.

troff allows the user full control over fonts, sizes, and character positions, as well as the usual features of a formatter — right-margin justification, automatic hyphenation, page titling and numbering, and so on. It also provides macros, arithmetic variables and operations, and conditional testing, for complicated formatting tasks.

This document is an introduction to the most basic use of **troff**. It presents just enough information to enable the user to do simple formatting tasks like making viewgraphs, and to make incremental changes to existing packages of **troff** commands. It assumes that the reader is familiar with a formatter like **roff** on UNIX or GCOS. In most respects, the UNIX formatter **nroff** is identical to **troff**, so this document also serves as a tutorial on **nroff**.



A TROFF Tutorial

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

I. Introduction

troff [1] is a text-formatting program, written by J. F. Ossanna, for producing high-quality printed output from the phototypesetter on the UNIX and GCOS operating systems. This document is an example of **troff** output.

The single most important rule of using **troff** is not to use it directly, but through some intermediary. In many ways, **troff** resembles an assembly language — a remarkably powerful and flexible one — but nonetheless such that many operations must be specified at a level of detail and in a form that is too hard for most people to use effectively.

For two special applications, there are programs that provide an interface to **troff** for the majority of users. **eqn** [2] provides an easy to learn language for typesetting mathematics; the **eqn** user need know no **troff** whatsoever to typeset mathematics. **tbl** [3] provides the same convenience for producing tables of arbitrary complexity.

For producing straight text (which may well contain mathematics or tables), there are a number of 'macro packages' that define formatting rules and operations for specific styles of documents, and reduce the amount of direct contact with **troff**. In particular, the '-ms' [4] and PWB/MM [5] packages for Bell Labs internal memoranda and external papers provide most of the facilities needed for a wide range of document preparation. (This memo was prepared with '-ms'.) There are also packages for viewgraphs, for simulating the older **roff** formatters on UNIX and GCOS, and for other special applications. Typically you will find these packages easier to use than **troff** once you get beyond the most trivial operations; you should always consider them first.

In the few cases where existing packages don't do the whole job, the solution is *not* to write an entirely new set of **troff** instructions from scratch, but to make small changes to adapt packages that already exist.

In accordance with this philosophy of letting someone else do the work, the part of **troff** described here is only a small part of the whole, although it tries to concentrate on the more useful parts. In any case, there is no attempt to be complete. Rather, the emphasis is on showing how to do simple things, and how to make incremental changes to what already exists. The contents of the remaining sections are:

2. Point sizes and line spacing
 3. Fonts and special characters
 4. Indents and line length
 5. Tabs
 6. Local motions: Drawing lines and characters
 7. Strings
 8. Introduction to macros
 9. Titles, pages and numbering
 10. Number registers and arithmetic
 11. Macros with arguments
 12. Conditionals
 13. Environments
 14. Diversions
- Appendix: Typesetter character set

The **troff** described here is the C-language version running on UNIX at Murray Hill, as documented in [1].

To use **troff** you have to prepare not only the actual text you want printed, but some information that tells *how* you want it printed. (Readers who use **roff** will find the approach familiar.) For **troff** the text and the formatting information are often intertwined quite intimately. Most commands to **troff** are placed on a line separate from the text itself, beginning with a period (one command per line). For example,

```
Some text.  
.ps 14  
Some more text.
```

will change the 'point size', that is, the size of the letters being printed, to '14 point' (one point is 1/72 inch) like this:

```
Some text. Some more text.
```

Occasionally, though, something special occurs in the middle of a line — to produce

$$\text{Area} = \pi r^2$$

you have to type

$$\text{Area} = \backslash(*p\flr\fr\|s8u2d\s0$$

(which we will explain shortly). The backslash character \ is used to introduce **troff** commands and special characters within a line of text.

2. Point Sizes; Line Spacing

As mentioned above, the command **.ps** sets the point size. One point is 1/72 inch, so 6-point characters are at most 1/12 inch high, and 36-point characters are 1/2 inch. There are 15 point sizes, listed below.

- 6 point: Pack my box with five dozen liquor jugs.
- 7 point: Pack my box with five dozen liquor jugs.
- 8 point: Pack my box with five dozen liquor jugs.
- 9 point: Pack my box with five dozen liquor jugs.
- 10 point: Pack my box with five dozen liquor
- 11 point: Pack my box with five dozen
- 12 point: Pack my box with five dozen
- 14 point: Pack my box with five
- 16 point
- 18 point
- 20 point
- 22
- 24
- 28
- 36

If the number after **.ps** is not one of these legal sizes, it is rounded up to the next valid value, with a maximum of 36. If no number follows **.ps**, **troff** reverts to the previous size, whatever it was. **troff** begins with point size 10, which is usually fine. This document is in 9 point.

The point size can also be changed in the middle of a line or even a word with the in-line command **\s**. To produce

UNIX runs on a PDP-11/45

type

\s8UNIX\s10 runs on a \s8PDP-\s1011/45

As above, **\s** should be followed by a legal point size, except that **\s0** causes the size to revert to its previous value. Notice that **\s1011** can be understood correctly as 'size 10, followed by an 11', if the size is legal, but not otherwise. Be cautious with similar constructions.

Relative size changes are also legal and useful:

\s-2UNIX\s+2

temporarily decreases the size, whatever it is, by two points, then restores it. Relative size changes have the advantage that the size difference is independent of the starting size of the document. The amount of the relative change is restricted to a single digit.

The other parameter that determines what the type looks like is the spacing between lines, which is set independently of the point size. Vertical spacing is measured from the bottom of one line to the bottom of the next. The command to control vertical spacing is **.vs**. For running text, it is usually best to set the vertical spacing about 20% bigger than the character size. For example, so far in this document, we have used "9 on 11", that is,

```
.ps 9
.vs 11p
```

If we changed to

```
.ps 9
.vs 9p
```

the running text would look like this. After a few lines, you will agree it looks a little cramped. The right vertical spacing is partly a matter of taste, depending on how much text you want to squeeze into a given space, and partly a matter of traditional printing style. By default, **troff** uses 10 on 12.

Point size and vertical spacing make a substantial difference in the amount of text per square inch. This is 12 on 14.

Point size and vertical spacing make a substantial difference in the amount of text per square inch. For example, 10 on 12 uses about twice as much space as 7 on 8. This is 6 on 7, which is even smaller. It packs a lot more words per line, but you can go blind trying to read it.

When used without arguments, **.ps** and **.vs** revert to the previous size and vertical spacing respectively.

The command **.sp** is used to get extra vertical space. Unadorned, it gives you one extra blank line (one **.vs**, whatever that has been set to). Typically, that's more or less than you want, so **.sp** can be followed by information about how much space you want —

```
.sp 2i
```

means 'two inches of vertical space'.

```
.sp 2p
```

means 'two points of vertical space'; and

.sp 2

means 'two vertical spaces' — two of whatever .vs is set to (this can also be made explicit with .sp 2v); troff also understands decimal fractions in most places, so

.sp 1.5i

is a space of 1.5 inches. These same scale factors can be used after .vs to define line spacing, and in fact after most commands that deal with physical dimensions.

It should be noted that all size numbers are converted internally to 'machine units', which are 1/432 inch (1/6 point). For most purposes, this is enough resolution that you don't have to worry about the accuracy of the representation. The situation is not quite so good vertically, where resolution is 1/144 inch (1/2 point).

3. Fonts and Special Characters

troff and the typesetter allow four different fonts at any one time. Normally three fonts (Times roman, italic and bold) and one collection of special characters are permanently mounted.

abcdefghijklmnopqrstuvwxy 0123456789
ABCDEFGHIJKLMNOPQRSTUVWXYZ
abcdefghijklmnopqrstuvwxy 0123456789
ABCDEFGHIJKLMNOPQRSTUVWXYZ
abcdefghijklmnopqrstuvwxy 0123456789
ABCDEFGHIJKLMNOPQRSTUVWXYZ

The greek, mathematical symbols and miscellany of the special font are listed in Appendix A.

troff prints in roman unless told otherwise. To switch into bold, use the .ft command

.ft B

and for italics,

.ft I

To return to roman, use .ft R; to return to the previous font, whatever it was, use either .ft P or just .ft. The 'underline' command

.ul

causes the next input line to print in italics. .ul can be followed by a count to indicate that more than one line is to be italicized.

Fonts can also be changed within a line or word with the in-line command \f:

bold*face* text

is produced by

\fBbold\fIface\fR text

If you want to do this so the previous font, whatever it was, is left undisturbed, insert extra \fP commands, like this:

\fBbold\fP\fIface\fP\fR text\fP

Because only the immediately previous font is remembered, you have to restore the previous font after each change or you can lose it. The same is true of .ps and .vs when used without an argument.

There are other fonts available besides the standard set, although you can still use only four at any given time. The command .fp tells troff what fonts are physically mounted on the typesetter:

.fp 3 H

says that the Helvetica font is mounted on position 3. (For a complete list of fonts and what they look like, see the troff manual.) Appropriate .fp commands should appear at the beginning of your document if you do not use the standard fonts.

It is possible to make a document relatively independent of the actual fonts used to print it by using font numbers instead of names; for example, \f3 and .ft 3 mean 'whatever font is mounted at position 3', and thus work for any setting. Normal settings are roman font on 1, italic on 2, bold on 3, and special on 4.

There is also a way to get 'synthetic' bold fonts by overstriking letters with a slight offset. Look at the .bd command in [1].

Special characters have four-character names beginning with \(), and they may be inserted anywhere. For example,

$\frac{1}{4} + \frac{1}{2} = \frac{3}{4}$

is produced by

\(14 + \(12 = \(34

In particular, greek letters are all of the form \(*—, where — is an upper or lower case roman letter reminiscent of the greek. Thus to get

$\Sigma(\alpha \times \beta) \rightarrow \infty$

in bare troff we have to type

\(*S\(*a\(\mu\(*b)\(-> \(\if

That line is unscrambled as follows:

<code>\(•S</code>	Σ
<code>(</code>	$($
<code>\(•a</code>	α
<code>\(mu</code>	\times
<code>\(•b</code>	β
<code>)</code>	$)$
<code>\(-></code>	\rightarrow
<code>\(if</code>	∞

A complete list of these special names occurs in Appendix A.

In `eqn` [2] the same effect can be achieved with the input

```
SIGMA ( alpha times beta ) -> inf
```

which is less concise, but clearer to the uninitiated.

Notice that each four-character name is a single character as far as `troff` is concerned — the ‘translate’ command

```
.tr \(\mi\)(em
```

is perfectly clear, meaning

```
.tr --
```

that is, to translate — into —.

Some characters are automatically translated into others: grave ` and acute ^ accents (apostrophes) become open and close single quotes ‘ ’; the combination of “...” is generally preferable to the double quotes "...". Similarly a typed minus sign becomes a hyphen -. To print an explicit — sign, use \-. To get a backslash printed, use \e.

4. Indents and Line Lengths

`troff` starts with a line length of 6.5 inches, too wide for 8½×11 paper. To reset the line length, use the `.ll` command, as in

```
.ll 6i
```

As with `.sp`, the actual length can be specified in several ways; inches are probably the most intuitive.

The maximum line length provided by the typesetter is 7.5 inches, by the way. To use the full width, you will have to reset the default physical left margin (“page offset”), which is normally slightly less than one inch from the left edge of the paper. This is done by the `.po` command.

```
.po 0
```

sets the offset as far to the left as it will go.

The indent command `.in` causes the left margin to be indented by some specified amount from the page offset. If we use `.in` to move the left margin in, and `.ll` to move the right margin to the left, we can make offset blocks of text:

```
.in 0.3i
.ll -0.3i
text to be set into a block
.ll +0.3i
.in -0.3i
```

will create a block that looks like this:

```
Pater noster qui est in caelis
sanctificetur nomen tuum; adveniat
regnum tuum; fiat voluntas tua, sicut
in caelo, et in terra. ... Amen.
```

Notice the use of ‘+’ and ‘-’ to specify the amount of change. These change the previous setting by the specified amount, rather than just overriding it. The distinction is quite important: `.ll +1i` makes lines one inch longer; `.ll 1i` makes them one inch *long*.

With `.in`, `.ll` and `.po`, the previous value is used if no argument is specified.

To indent a single line, use the ‘temporary indent’ command `.ti`. For example, all paragraphs in this memo effectively begin with the command

```
.ti 3
```

Three of what? The default unit for `.ti`, as for most horizontally oriented commands (`.ll`, `.in`, `.po`), is ems; an em is roughly the width of the letter ‘m’ in the current point size. (Precisely, a em in size *p* is *p* points.) Although inches are usually clearer than ems to people who don’t set type for a living, ems have a place: they are a measure of size ‘that is proportional to the current point size. If you want to make text that keeps its proportions regardless of point size, you should use ems for all dimensions. Ems can be specified as scale factors directly, as in `.ti 2.5m`.

Lines can also be indented negatively if the indent is already positive:

```
.ti -0.3i
```

causes the next line to be moved back three tenths of an inch. Thus to make a decorative initial capital, we indent the whole paragraph, then move the letter ‘P’ back with a `.ti` command:

Pater noster qui est in caelis
 sanctificetur nomen tuum; ad-
 veniat regnum tuum; fiat volun-
 tas tua, sicut in caelo, et in terra. ...
 Amen.

Of course, there is also some trickery to make the 'P' bigger (just a '\s36P\s0'), and to move it down from its normal position (see the section on local motions).

5. Tabs

Tabs (the ASCII 'horizontal tab' character) can be used to produce output in columns, or to set the horizontal position of output. Typically tabs are used only in unfilled text. Tab stops are set by default every half inch from the current indent, but can be changed by the .ta command. To set stops every inch, for example,

```
.ta 1i 2i 3i 4i 5i 6i
```

Unfortunately the stops are left-justified only (as on a typewriter), so lining up columns of right-justified numbers can be painful. If you have many numbers, or if you need more complicated table layout, *don't* use **troff** directly; use the **tbl** program described in [3].

For a handful of numeric columns, you can do it this way: Precede every number by enough blanks to make it line up when typed.

```
.nf
.ta 1i 2i 3i
  1 tab 2 tab 3
 40 tab 50 tab 60
700 tab 800 tab 900
.fi
```

Then change each leading blank into the string \0. This is a character that does not print, but that has the same width as a digit. When printed, this will produce

1	2	3
40	50	60
700	800	900

It is also possible to fill up tabbed-over space with some character other than blanks by setting the 'tab replacement character' with the .tc command:

```
.ta 1.5i 2.5i
.tc \ (ru (\ (ru is "-")
Name tab Age tab
```

produces

Name _____ Age _____

To reset the tab replacement character to a blank, use .tc with no argument. (Lines can also be drawn with the \l command, described in Section 6.)

troff also provides a very general mechanism called 'fields' for setting up complicated columns. (This is used by **tbl**). We will not go into it in this paper.

6. Local Motions: Drawing lines and characters

Remember 'Area = πr^2 ', and the big 'P' in the Paternoster. How are they done? **troff** provides a host of commands for placing characters of any size at any place. You can use them to draw special characters or to tune your output for a particular appearance. Most of these commands are straightforward, but messy to read and tough to type correctly.

If you won't use **eqn**, subscripts and superscripts are most easily done with the half-line local motions \u and \d. To go back up the page half a point-size, insert a \u at the desired place; to go down, insert a \d. (\u and \d should always be used in pairs, as explained below.) Thus

```
Area = \(*pr\u2\d
```

produces

$$\text{Area} = \pi r^2$$

To make the '2' smaller, bracket it with \s-2..\s0. Since \u and \d refer to the current point size, be sure to put them either both inside or both outside the size changes, or you will get an unbalanced vertical motion.

Sometimes the space given by \u and \d isn't the right amount. The \v command can be used to request an arbitrary amount of vertical motion. The in-line command

```
\v'(amount)'
```

causes motion up or down the page by the amount specified in '(amount)'. For example, to move the 'P' down, we used

```
.in +0.6i      (move paragraph in)
.ll -0.3i     (shorten lines)
.ti -0.3i     (move P back)
\v'2'\s36P\s0\v'-2'ater noster qui est
in caelis ...
```

A minus sign causes upward motion, while no sign or a plus sign means down the page. Thus \v'-2' causes an upward vertical motion of two line spaces.

There are many other ways to specify the amount of motion —

```
\v'0.1i'
\v'3p'
\v'-0.5m'
```

and so on are all legal. Notice that the scale specifier *i* or *p* or *m* goes inside the quotes. Any character can be used in place of the quotes; this is also true of all other **troff** commands described in this section.

Since **troff** does not take within-the-line vertical motions into account when figuring out where it is on the page, output lines can have unexpected positions if the left and right ends aren't at the same vertical position. Thus $\backslash v$, like $\backslash u$ and $\backslash d$, should always balance upward vertical motion in a line with the same amount in the downward direction.

Arbitrary horizontal motions are also available — $\backslash h$ is quite analogous to $\backslash v$, except that the default scale factor is ems instead of line spaces. As an example,

```
\h'-0.1i'
```

causes a backwards motion of a tenth of an inch. As a practical matter, consider printing the mathematical symbol '>>'. The default spacing is too wide, so **eqn** replaces this by

```
>\h'-0.3m'>
```

to produce >>.

Frequently $\backslash h$ is used with the 'width function' $\backslash w$ to generate motions equal to the width of some character string. The construction

```
\w'thing'
```

is a number equal to the width of 'thing' in machine units (1/432 inch). All **troff** computations are ultimately done in these units. To move horizontally the width of an 'x', we can say

```
\h'\w'x'u'
```

As we mentioned above, the default scale factor for all horizontal dimensions is *m*, ems, so here we must have the *u* for machine units, or the motion produced will be far too large. **troff** is quite happy with the nested quotes, by the way, so long as you don't leave any out.

As a live example of this kind of construction, all of the command names in the text, like **.sp**, were done by overstriking with a slight offset. The commands for **.sp** are

```
.sp\h'-\w'.sp'u'\h'1u'.sp
```

That is, put out **.sp'**, move left by the width of **.sp'**, move right 1 unit, and print **.sp'** again. (Of course there is a way to avoid typing that much input for each command name, which we will discuss in Section 11.)

There are also several special-purpose **troff** commands for local motion. We have already seen $\backslash 0$, which is an unpaddable white space of the same width as a digit. 'Unpaddable' means that it will never be widened or split across a line by line justification and filling. There is also $\backslash (\text{blank})$, which is an unpaddable character the width of a space, $\backslash \downarrow$, which is half that width, $\backslash \wedge$, which is one quarter of the width of a space, and $\backslash \&$, which has zero width. (This last one is useful, for example, in entering a text line which would otherwise begin with a '.')

The command $\backslash o$, used like

```
\o'set of characters'
```

causes (up to 9) characters to be overstruck, centered on the widest. This is nice for accents, as in

```
syst\o"e\(\ga"me t\o"e\(\aa"l\o"e\(\aa"phonique
```

which makes

```
systeme téléphonique
```

The accents are $\backslash (\text{ga})$ and $\backslash (\text{aa})$, or $\backslash \grave$ and $\backslash \text{;}$; remember that each is just one character to **troff**.

You can make your own overstrikes with another special convention, $\backslash z$, the zero-motion command. $\backslash zx$ suppresses the normal horizontal motion after printing the single character *x*, so another character can be laid on top of it. Although sizes can be changed within $\backslash o$, it centers the characters on the widest, and there can be no horizontal or vertical motions, so $\backslash z$ may be the only way to get what you want:



is produced by

```
.sp 2
\s8\z\(\sq\s14\z\(\sq\s22\z\(\sq\s36\(\sq
```

The **.sp** is needed to leave room for the result.

As another example, an extra-heavy semicolon that looks like

```
;
```

can be constructed with a big comma and a big period above it:

`\s + 6\z,\v' - 0.25m' \v'0.25m'\s0`

'0.25m' is an empirical constant.

A more ornate overstrike is given by the bracketing function `\b`, which piles up characters vertically, centered on the current baseline. Thus we can get big brackets, constructing them with piled-up smaller pieces:

$$\left\{ \left[x \right] \right\}$$

by typing in only this:

```
.sp
\b'\(l\)(k\)(l'b'\(l\)(l' x \b'\(r\)(r' \b'\(r\)(r'
```

troff also provides a convenient facility for drawing horizontal and vertical lines of arbitrary length with arbitrary characters. `\l'i` draws a line one inch long, like this: _____ . The length can be followed by the character to use if the `_` isn't appropriate; `\l'0.5i.'` draws a half-inch line of dots: The construction `\L` is entirely analogous, except that it draws a vertical line instead of horizontal.

7. Strings

Obviously if a paper contains a large number of occurrences of an acute accent over a letter 'e', typing `\o"e\"` for each é would be a great nuisance.

Fortunately, **troff** provides a way in which you can store an arbitrary collection of text in a 'string', and thereafter use the string name as a shorthand for its contents. Strings are one of several **troff** mechanisms whose judicious use lets you type a document with less effort and organize it so that extensive format changes can be made with few editing changes.

A reference to a string is replaced by whatever text the string was defined as. Strings are defined with the command `.ds`. The line

```
.ds e \o"e\"
```

defines the string `e` to have the value `\o"e\"`

String names may be either one or two characters long, and are referred to by `*x` for one character names or `*(xy` for two character names. Thus to get `téléphone`, given the definition of the string `e` as above, we can say `t*el*ephone`.

If a string must begin with blanks, define it as

```
.ds xx " text
```

The double quote signals the beginning of the definition. There is no trailing quote; the end of the line terminates the string.

A string may actually be several lines long; if **troff** encounters a `\` at the end of *any* line, it is thrown away and the next line added to the current one. So you can make a long string simply by ending each line but the last with a backslash:

```
.ds xx this \
is a very \
long string
```

Strings may be defined in terms of other strings, or even in terms of themselves; we will discuss some of these possibilities later.

8. Introduction to Macros

Before we can go much further in **troff**, we need to learn a bit about the macro facility. In its simplest form, a macro is just a shorthand notation quite similar to a string. Suppose we want every paragraph to start in exactly the same way — with a space and a temporary indent of two ems:

```
.sp
.ti +2m
```

Then to save typing, we would like to collapse these into one shorthand line, a **troff** 'command' like

```
.PP
```

that would be treated by **troff** exactly as

```
.sp
.ti +2m
```

`.PP` is called a *macro*. The way we tell **troff** what `.PP` means is to *define* it with the `.de` command:

```
.de PP
.sp
.ti +2m
..
```

The first line names the macro (we used `.PP` for 'paragraph', and upper case so it wouldn't conflict with any name that **troff** might already know about). The last line `..` marks the end of the definition. In between is the text, which is simply inserted whenever **troff** sees the 'command' or macro call

```
.PP
```

A macro can contain any mixture of text and formatting commands.

The definition of `.PP` has to precede its first use; undefined macros are simply ignored. Names are restricted to one or two characters.

Using macros for commonly occurring sequences of commands is critically important. **Not only** does it save typing, but it makes later changes much easier. Suppose we decide that the paragraph indent is too small, the vertical space is much too big, and roman font should be forced. Instead of changing the whole document, we need only change the definition of `.PP` to something like

```
.de PP      \" paragraph macro
.sp 2p
.tl +3m
.ft R
..
```

and the change takes effect everywhere we used `.PP`.

`\"` is a **troff** command that causes the rest of the line to be ignored. We use it here to add comments to the macro definition (a wise idea once definitions get complicated).

As another example of macros, consider these two which start and end a block of offset, unfilled text, like most of the examples in this paper:

```
.de BS      \" start indented block
.sp
.nf
.in +0.3i
..
.de BE      \" end indented block
.sp
.fi
.in -0.3i
..
```

Now we can surround text like

```
Copy to
John Doe
Richard Roberts
Stanley Smith
```

by the commands `.BS` and `.BE`, and it will come out as it did above. Notice that we indented by `.in +0.3i` instead of `.in 0.3i`. This way we can nest our uses of `.BS` and `.BE` to get blocks within blocks.

If later on we decide that the indent should be `0.5i`, then it is only necessary to change the definitions of `.BS` and `.BE`, not the whole paper.

9. Titles, Pages and Numbering

This is an area where things get tougher, because nothing is done for you automatically. Of necessity, some of this section is a cookbook, to be copied literally until you get some experience.

Suppose you want a title at the top of each page, saying just

```
left top          center top          right top
```

In **roff**, one can say

```
.he 'left top'center top'right top'
.fh 'left bottom'center bottom'right bottom'
```

to get headers and footers automatically on every page. Alas, this doesn't work in **troff**, a serious hardship for the novice. Instead you have to do a lot of specification.

You have to say what the actual title is (easy); when to print it (easy enough); and what to do at and around the title line (harder). Taking these in reverse order, first we define a macro `.NP` (for 'new page') to process titles and the like at the end of one page and the beginning of the next:

```
.de NP
'bp
'sp 0.5i
.tl 'left top'center top'right top'
'sp 0.3i
..
```

To make sure we're at the top of a page, we issue a 'begin page' command `'bp`, which causes a skip to top-of-page (we'll explain the ' shortly). Then we space down half an inch, print the title (the use of `.tl` should be self explanatory; later we will discuss parameterizing the titles), space another 0.3 inches, and we're done.

To ask for `.NP` at the bottom of each page, we have to say something like 'when the text is within an inch of the bottom of the page, start the processing for a new page.' This is done with a 'when' command `.wh`:

```
.wh -1i NP
```

(No '.' is used before `NP`; this is simply the name of a macro, not a macro call.) The minus sign means 'measure up from the bottom of the page', so `'-1i'` means 'one inch from the bottom'.

The `.wh` command appears in the input outside the definition of `.NP`; typically the input would be

```
.de NP
...
...
.wh -li NP
```

Now what happens? As text is actually being output, **troff** keeps track of its vertical position on the page, and after a line is printed within one inch from the bottom, the **.NP** macro is activated. (In the jargon, the **.wh** command sets a *trap* at the specified place, which is 'sprung' when that point is passed.) **.NP** causes a skip to the top of the next page (that's what the **'bp** was for), then prints the title with the appropriate margins.

Why **'bp** and **'sp** instead of **.bp** and **.sp**? The answer is that **.sp** and **.bp**, like several other commands, cause a *break* to take place. That is, all the input text collected but not yet printed is flushed out as soon as possible, and the next input line is guaranteed to start a new line of output. If we had used **.sp** or **.bp** in the **.NP** macro, this would cause a break in the middle of the current output line when a new page is started. The effect would be to print the left-over part of that line at the top of the page, followed by the next input line on a new output line. This is *not* what we want. Using **'** instead of **.** for a command tells **troff** that no break is to take place — the output line currently being filled should *not* be forced out before the space or new page.

The list of commands that cause a break is short and natural:

```
.bp .br .ce .fi .nf .sp .in .ti
```

All others cause *no* break, regardless of whether you use a **.** or a **'**. If you really need a break, add a **.br** command at the appropriate place.

One other thing to beware of — if you're changing fonts or point sizes a lot, you may find that if you cross a page boundary in an unexpected font or size, your titles come out in that size and font instead of what you intended. Furthermore, the length of a title is independent of the current line length, so titles will come out at the default length of 6.5 inches unless you change it, which is done with the **.lt** command.

There are several ways to fix the problems of point sizes and fonts in titles. For the simplest applications, we can change **.NP** to set the proper size and font for the title, then restore the previous values, like this:

```
.de NP
'bp
'sp 0.5i
.ft R          \" set title font to roman
.ps 10         \" and size to 10 point
.lt 6i        \" and length to 6 inches
.tl 'left'center'right'
.ps           \" revert to previous size
.ft P         \" and to previous font
'sp 0.3i
..
```

This version of **.NP** does *not* work if the fields in the **.tl** command contain size or font changes. To cope with that requires **troff**'s 'environment' mechanism, which we will discuss in Section 13.

To get a footer at the bottom of a page, you can modify **.NP** so it does some processing before the **'bp** command, or split the job into a footer macro invoked at the bottom margin and a header macro invoked at the top of the page. These variations are left as exercises.

Output page numbers are computed automatically as each page is produced (starting at 1), but no numbers are printed unless you ask for them explicitly. To get page numbers printed, include the character **%** in the **.tl** line at the position where you want the number to appear. For example

```
.tl "- % -"
```

centers the page number inside hyphens, as on this page. You can set the page number at any time with either **.bp n**, which immediately starts a new page numbered **n**, or with **.pn n**, which sets the page number for the next page but doesn't cause a skip to the new page. Again, **.bp +n** sets the page number to **n** more than its current value; **.bp** means **.bp +1**.

10. Number Registers and Arithmetic

troff has a facility for doing arithmetic, and for defining and using variables with numeric values, called *number registers*. Number registers, like strings and macros, can be useful in setting up a document so it is easy to change later. And of course they serve for any sort of arithmetic computation.

Like strings, number registers have one or two character names. They are set by the **.nr** command, and are referenced anywhere by **\nx** (one character name) or **\n(xy** (two character name).

There are quite a few pre-defined number registers maintained by **troff**, among them **%** for the current page number; **nl** for the current vertical position on the page; **dy**, **mo** and **yr** for the current day, month and year; and **.s** and **.f** for the current size and font. (The font is a number from 1 to 4.) Any of these can be used in computations like any other register, but some, like **.s** and **.f**, cannot be changed with **.nr**.

As an example of the use of number registers, in the **-ms** macro package [4], most significant parameters are defined in terms of the values of a handful of number registers. These include the point size for text, the vertical spacing, and the line and title lengths. To set the point size and vertical spacing for the following paragraphs, for example, a user may say

```
.nr PS 9
.nr VS 11
```

The paragraph macro **.PP** is defined (roughly) as follows:

```
.de PP
.ps \n(PS      \" reset size
.vs \n(VSp     \" spacing
.ft R         \" font
.sp 0.5v      \" half a line
.ti +3m
..
```

This sets the font to Roman and the point size and line spacing to whatever values are stored in the number registers **PS** and **VS**.

Why are there two backslashes? This is the eternal problem of how to quote a quote. When **troff** originally reads the macro definition, it peels off one backslash to see what's coming next. To ensure that another is left in the definition when the macro is used, we have to put in two backslashes in the definition. If only one backslash is used, point size and vertical spacing will be frozen at the time the macro is defined, not when it is used.

Protecting by an extra layer of backslashes is only needed for **\n**, *****, **\\$** (which we haven't come to yet), and **** itself. Things like **\s**, **\f**, **\h**, **\v**, and so on do not need an extra backslash, since they are converted by **troff** to an internal code immediately upon being seen.

Arithmetic expressions can appear anywhere that a number is expected. As a trivial example,

```
.nr PS \n(PS-2
```

decrements **PS** by 2. Expressions can use the

arithmetic operators **+**, **-**, *****, **/**, **%** (mod), the relational operators **>**, **>=**, **<**, **<=**, **=**, and **!=** (not equal), and parentheses.

Although the arithmetic we have done so far has been straightforward, more complicated things are somewhat tricky. First, number registers hold only integers. **troff** arithmetic uses truncating integer division, just like Fortran. Second, in the absence of parentheses, evaluation is done left-to-right without any operator precedence (including relational operators). Thus

```
7*-4+3/13
```

becomes **'-1'**. Number registers can occur anywhere in an expression, and so can scale indicators like **p**, **i**, **m**, and so on (but no spaces). Although integer division causes truncation, each number and its scale indicator is converted to machine units (1/432 inch) before any arithmetic is done, so **1i/2u** evaluates to **0.5i** correctly.

The scale indicator **u** often has to appear when you wouldn't expect it — in particular, when arithmetic is being done in a context that implies horizontal or vertical dimensions. For example,

```
.ll 7/2i
```

would seem obvious enough — **3½** inches. Sorry. Remember that the default units for horizontal parameters like **.ll** are ems. That's really **'7 ems / 2 inches'**, and when translated into machine units, it becomes zero. How about

```
.ll 7i/2
```

Sorry, still no good — the **'2'** is **'2 ems'**, so **'7i/2'** is small, although not zero. You *must* use

```
.ll 7i/2u
```

So again, a safe rule is to attach a scale indicator to every number, even constants.

For arithmetic done within a **.nr** command, there is no implication of horizontal or vertical dimension, so the default units are **'units'**, and **7i/2** and **7i/2u** mean the same thing. Thus

```
.nr ll 7i/2
.ll \n(llu
```

does just what you want, so long as you don't forget the **u** on the **.ll** command.

11. Macros with arguments

The next step is to define macros that can change from one use to the next according to parameters supplied as arguments. To make this work, we need two things: first, when we define

the macro, we have to indicate that some parts of it will be provided as arguments when the macro is called. Then when the macro is called we have to provide actual arguments to be plugged into the definition.

Let us illustrate by defining a macro **.SM** that will print its argument two points smaller than the surrounding text. That is, the macro call

```
.SM TROFF
```

will produce TROFF.

The definition of **.SM** is

```
.de SM
\s-2\\$1\s+2
```

Within a macro definition, the symbol **\\\$n** refers to the *n*th argument that the macro was called with. Thus **\\\$1** is the string to be placed in a smaller point size when **.SM** is called.

As a slightly more complicated version, the following definition of **.SM** permits optional second and third arguments that will be printed in the normal size:

```
.de SM
\\$3\s-2\\$1\s+2\\$2
```

Arguments not provided when the macro is called are treated as empty, so

```
.SM TROFF ),
```

produces TROFF), while

```
.SM TROFF ). (
```

produces (TROFF). It is convenient to reverse the order of arguments because trailing punctuation is much more common than leading.

By the way, the number of arguments that a macro was called with is available in number register **.S**.

The following macro **.BD** is the one used to make the 'bold roman' we have been using for **troff** command names in text. It combines horizontal motions, width computations, and argument rearrangement.

```
.de BD
\\&\\$3\\f1\\$1\\b'-\\w\\$1'u+1u\\$1\\fP\\$2
```

The **\\h** and **\\w** commands need no extra backslash, as we discussed above. The **\\&** is there in case the argument begins with a period.

Two backslashes are needed with the **\\\$n** commands, though, to protect one of them when the macro is being defined. Perhaps a second example will make this clearer. Consider a macro called **.SH** which produces section headings rather like those in this paper, with the sections numbered automatically, and the title in bold in a smaller size. The use is

```
.SH "Section title ..."
```

(If the argument to a macro is to contain blanks, then it must be *surrounded* by double quotes, unlike a string, where only one leading quote is permitted.)

Here is the definition of the **.SH** macro:

```
.nr SH 0      \" initialize section number
.de SH
.sp 0.3i
.ft B
.nr SH \\n(SH+1  \" increment number
.ps \\n(PS-1     \" decrease PS
\\n(SH. \\$1     \" number, title
.ps \\n(PS       \" restore PS
.sp 0.3i
.ft R
```

The section number is kept in number register **SH**, which is incremented each time just before it is used. (A number register may have the same name as a macro without conflict but a string may not.)

We used **\\n(SH** instead of **\\n(SH** and **\\n(PS** instead of **\\n(PS**. If we had used **\\n(SH**, we would get the value of the register at the time the macro was *defined*, not at the time it was *used*. If that's what you want, fine, but not here. Similarly, by using **\\n(PS**, we get the point size at the time the macro is called.

As an example that does not involve numbers, recall our **.NP** macro which had a

```
.tl 'left'center'right'
```

We could make these into parameters by using instead

```
.tl '\\*(LT'\\*(CT'\\*(RT'
```

so the title comes from three strings called **LT**, **CT** and **RT**. If these are empty, then the title will be a blank line. Normally **CT** would be set with something like

```
.ds CT - % -
```

to give just the page number between hyphens (as on the top of this page), but a user could

supply private definitions for any of the strings.

12. Conditionals

Suppose we want the `.SH` macro to leave two extra inches of space just before section 1, but nowhere else. The cleanest way to do that is to test inside the `.SH` macro whether the section number is 1, and add some space if it is. The `.if` command provides the conditional test that we can add just before the heading line is output:

```
.if \n(SH=1 .sp 2i    \" first section only
```

The condition after the `.if` can be any arithmetic or logical expression. If the condition is logically true, or arithmetically greater than zero, the rest of the line is treated as if it were text → here a command. If the condition is false, or zero or negative, the rest of the line is skipped.

It is possible to do more than one command if a condition is true. Suppose several operations are to be done before section 1. One possibility is to define a macro `.S1` and invoke it if we are about to do section 1 (as determined by an `.if`).

```
.de S1
--- processing for section 1 ---
...
.de SH
...
.if \n(SH=1 .S1
...
.
```

An alternate way is to use the extended form of the `.if`, like this:

```
.if \n(SH=1 \{--- processing
for section 1 ---\}
```

The braces `{` and `}` must occur in the positions shown or you will get unexpected extra lines in your output. `troff` also provides an 'if-else' construction, which we will not go into here.

A condition can be negated by preceding it with `!`; we get the same effect as above (but less clearly) by using

```
.if !\n(SH>1 .S1
```

There are a handful of other conditions that can be tested with `.if`. For example, is the current page even or odd?

```
.if e .tl "even page title"
.if o .tl "odd page title"
```

gives facing pages different titles when used

inside an appropriate new page macro.

Two other conditions are `t` and `n`, which tell you whether the formatter is `troff` or `nroff`.

```
.if t troff stuff ...
.if n nroff stuff ...
```

Finally, string comparisons may be made in an `.if`:

```
.if 'string1'string2' stuff
```

does 'stuff' if *string1* is the same as *string2*. The character separating the strings can be anything reasonable that is not contained in either string. The strings themselves can reference strings with `*`, arguments with `\$`, and so on.

13. Environments

As we mentioned, there is a potential problem when going across a page boundary: parameters like size and font for a page title may well be different from those in effect in the text when the page boundary occurs. `troff` provides a very general way to deal with this and similar situations. There are three 'environments', each of which has independently settable versions of many of the parameters associated with processing, including size, font, line and title lengths, fill/nofill mode, tab stops, and even partially collected lines. Thus the titling problem may be readily solved by processing the main text in one environment and titles in a separate one with its own suitable parameters.

The command `.ev n` shifts to environment `n`; `n` must be 0, 1 or 2. The command `.ev` with no argument returns to the previous environment. Environment names are maintained in a stack, so calls for different environments may be nested and unwound consistently.

Suppose we say that the main text is processed in environment 0, which is where `troff` begins by default. Then we can modify the new page macro `.NP` to process titles in environment 1 like this:

```
.de NP
.ev 1    \" shift to new environment
.lt 6i   \" set parameters here
.ft R
.ps 10
... any other processing ...
.ev     \" return to previous environment
.
```

It is also possible to initialize the parameters for an environment outside the `.NP` macro, but the version shown keeps all the processing in one place and is thus easier to understand and

change.

14. Diversions

There are numerous occasions in page layout when it is necessary to store some text for a period of time without actually printing it. Footnotes are the most obvious example: the text of the footnote usually appears in the input well before the place on the page where it is to be printed is reached. In fact, the place where it is output normally depends on how big it is, which implies that there must be a way to process the footnote at least enough to decide its size without printing it.

troff provides a mechanism called a diversion for doing this processing. Any part of the output may be diverted into a macro instead of being printed, and then at some convenient time the macro may be put back into the input.

The command `.di xy` begins a diversion — all subsequent output is collected into the macro `xy` until the command `.di` with no arguments is encountered. This terminates the diversion. The processed text is available at any time thereafter, simply by giving the command

```
.xy
```

The vertical size of the last finished diversion is contained in the built-in number register `dn`.

As a simple example, suppose we want to implement a 'keep-release' operation, so that text between the commands `.KS` and `.KE` will not be split across a page boundary (as for a figure or table). Clearly, when a `.KS` is encountered, we have to begin diverting the output so we can find out how big it is. Then when a `.KE` is seen, we decide whether the diverted text will fit on the current page, and print it either there if it fits, or at the top of the next page if it doesn't. So:

```

.de KS    \" start keep
.br      \" start fresh line
.ev l    \" collect in new environment
.fi      \" make it filled text
.di XX   \" collect in XX
...

.de KE    \" end keep
.br      \" get last partial line
.di      \" end diversion
.if \\n(dn >= \\n(.t.bp) \" bp if doesn't fit
.nf      \" bring it back in no-fill
.XX      \" text
.ev      \" return to normal environment
...

```

Recall that number register `nl` is the current

position on the output page. Since output was being diverted, this remains at its value when the diversion started. `dn` is the amount of text in the diversion; `.t` (another built-in register) is the distance to the next trap, which we assume is at the bottom margin of the page. If the diversion is large enough to go past the trap, the `.if` is satisfied, and a `.bp` is issued. In either case, the diverted output is then brought back with `.XX`. It is essential to bring it back in no-fill mode so **troff** will do no further processing on it.

This is not the most general keep-release, nor is it robust in the face of all conceivable inputs, but it would require more space than we have here to write it in full generality. This section is not intended to teach everything about diversions, but to sketch out enough that you can read existing macro packages with some comprehension.

Acknowledgements

I am deeply indebted to J. F. Ossanna, the author of **troff**, for his repeated patient explanations of fine points, and for his continuing willingness to adapt **troff** to make other uses easier. I am also grateful to Jim Blinn, Ted Dolotta, Doug McIlroy, Mike Lesk and Joel Sturman for helpful comments on this paper.

References

- [1] J. F. Ossanna, *NROFF/TROFF User's Manual*, Bell Laboratories internal memorandum.
- [2] B. W. Kernighan, *A System for Typesetting Mathematics — User's Guide (Second Edition)*, Bell Laboratories internal memorandum.
- [3] M. E. Lesk, *TBL — A Program to Format Tables*, Bell Laboratories internal memorandum.
- [4] M. E. Lesk, *Typing Documents on UNIX*, Bell Laboratories internal memorandum.
- [5] J. R. Mashey and D. W. Smith, *PWB/IMM — Programmer's Workbench Memorandum Macros*, Bell Laboratories internal memorandum.

Appendix A: Phototypesetter Character Set

These characters exist in roman, italic, and bold. To get the one on the left, type the four-character name on the right.

fl	\(fl	li	\(li	fl	\(fl	fli	\(Fi	fl	\(Fl
-	\(ru	-	\(em	¼	\(14	½	\(12	¼	\(34
°	\(co	°	\(de	†	\(dg	'	\(fm	¢	\(ct
•	\(rg	•	\(bu	□	\(sq	-	\(hy		

(In bold, \(\sq is ■.)

The following are special-font characters:

+	\(pl	-	\(mi	×	\(mu	÷	\(di
=	\(eq	≡	\(==	≥	\(>=	≤	\(<=
≠	\(!=	±	\(+-	¬	\(no	/	\(sl
~	\(ap	≈	\(≈	∞	\(pt	∇	\(gr
→	\(>	←	\(<	↑	\(ua	↓	\(da
∫	\(is	∂	\(pd	∞	\(if	√	\(sr
⊂	\(sb	⊃	\(sp	∪	\(cu	∩	\(ca
⊆	\(ib	⊇	\(ip	∈	\(mo	∅	\(es
'	\(aa	'	\(ga	○	\(ci	⊕	\(bs
§	\(sc	‡	\(dd	■	\(lh	■	\(rh
	\(lt		\(rt		\(lc		\(rc
	\(lb		\(rb		\(lf		\(rf
	\(lk		\(rk		\(bv	ˆ	\(ts
	\(br		\(or	-	\(ul	-	\(rn
•	\(**						

These four characters also have two-character names. The ´ is the apostrophe on terminals; the ` is the other quote mark.

´	\('	´	\('	-	\(-	-	\(-
---	-----	---	-----	---	-----	---	-----


These characters exist only on the special font, but they do not have four-character names:

"	{	}	<	>	ˆ	ˆ	\	#	@
---	---	---	---	---	---	---	---	---	---

For greek, precede the roman letter by \(* to get the corresponding greek; for example, \(*a is α.

a	b	g	d	e	z	y	h	i	k	l	m	n	c	o	p	r	s	t	u	f	x	q	w
α	β	γ	δ	ε	ζ	η	θ	ι	κ	λ	μ	ν	ξ	ο	π	ρ	σ	τ	υ	φ	χ	ψ	ω
A	B	G	D	E	Z	Y	H	I	K	L	M	N	C	O	P	R	S	T	U	F	X	Q	W
Α	Β	Γ	Δ	Ε	Ζ	Η	Θ	Ι	Κ	Λ	Μ	Ν	Ξ	Ο	Π	Ρ	Σ	Τ	Υ	Φ	Χ	Ψ	Ω





Typesetting Mathematics — User's Guide (Second Edition)

Brian W. Kernighan

Bell Laboratories,
Murray Hill, New Jersey 07974

Lorinda L. Cherry

Bell Laboratories,
Murray Hill, New Jersey 07974

ABSTRACT

This is the user's guide for a system for typesetting mathematics, using the phototypesetters on the UNIX and GCOS operating systems.

Mathematical expressions are described in a language designed to be easy to use by people who know neither mathematics nor typesetting. Enough of the language to set in-line expressions like $\lim_{x \rightarrow \pi/2} (\tan x)^{\sin 2x} = 1$ or display equations like

$$\begin{aligned}
 G(z) &= e^{\ln G(z)} = \exp \left(\sum_{k \geq 1} \frac{S_k z^k}{k} \right) = \prod_{k \geq 1} e^{S_k z^k / k} \\
 &= \left(1 + S_1 z + \frac{S_1^2 z^2}{2!} + \dots \right) \left(1 + \frac{S_2 z^2}{2} + \frac{S_2^2 z^4}{2^2 \cdot 2!} + \dots \right) \dots \\
 &= \sum_{m \geq 0} \left(\sum_{\substack{k_1, k_2, \dots, k_m \geq 0 \\ k_1 + 2k_2 + \dots + mk_m = m}} \frac{S_1^{k_1}}{1^{k_1} k_1!} \frac{S_2^{k_2}}{2^{k_2} k_2!} \dots \frac{S_m^{k_m}}{m^{k_m} k_m!} \right) z^m
 \end{aligned}$$

can be learned in an hour or so.

The language interfaces directly with the phototypesetting language TROFF, so mathematical expressions can be embedded in the running text of a manuscript, and the entire document produced in one process. This user's guide is an example of its output.

The same language may be used with the UNIX formatter NROFF to set mathematical expressions on DASI and GSI terminals and Model 37 teletypes.



Typesetting Mathematics — User's Guide (Second Edition)

Brian W. Kernighan

Bell Laboratories,
Murray Hill, New Jersey 07974

Lorinda L. Cherry

Bell Laboratories,
Murray Hill, New Jersey 07974

1. Introduction

EQN is a program for typesetting mathematics on the Graphics Systems phototypesetters on UNIX and GCOS. The EQN language was designed to be easy to use by people who know neither mathematics nor typesetting. Thus EQN knows relatively little about mathematics. In particular, mathematical symbols like +, −, ×, parentheses, and so on have no special meanings. EQN is quite happy to set garbage (but it will look good).

EQN works as a preprocessor for the typesetter formatter, TROFF[1], so the normal mode of operation is to prepare a document with both mathematics and ordinary text interspersed, and let EQN set the mathematics while TROFF does the body of the text.

On UNIX, EQN will also produce mathematics on DASI and GSI terminals and on Model 37 teletypes. The input is identical, but you have to use the programs NEQN and NROFF instead of EQN and TROFF. Of course, some things won't look as good because terminals don't provide the variety of characters, sizes and fonts that a typesetter does, but the output is usually adequate for proofreading.

To use EQN on UNIX,

```
eqn files | troff
```

GCOS use is discussed in section 26.

2. Displayed Equations

To tell EQN where a mathematical expression begins and ends, we mark it with lines beginning EQ and EN. Thus if you type the lines

```
.EQ  
x=y+z  
.EN
```

your output will look like

$$x=y+z$$

The EQ and EN are copied through untouched; they are not otherwise processed by EQN. This means that you have to take care of things like centering, numbering, and so on yourself. The most common way is to use the TROFF and NROFF macro package package '-ms' developed by M. E. Lesk[3], which allows you to center, indent, left-justify and number equations.

With the '-ms' package, equations are centered by default. To left-justify an equation, use EQ L instead of EQ. To indent it, use EQ I. Any of these can be followed by an arbitrary 'equation number' which will be placed at the right margin. For example, the input

```
.EQ I (3.1a)  
x = f(y/2) + y/2  
.EN
```

produces the output

$$x=f(y/2)+y/2 \qquad (3.1a)$$

There is also a shorthand notation so in-line expressions like π^2 can be entered without EQ and EN. We will talk about it in section 19.

3. Input spaces

Spaces and newlines within an expression are thrown away by EQN. (Normal text is left absolutely alone.) Thus between EQ and EN,

$$x=y+z$$

and

$$x = y + z$$

and

$$x = y \\ + z$$

and so on all produce the same output

$$x=y+z$$

You should use spaces and newlines freely to make your input equations readable and easy to edit. In particular, very long lines are a bad idea, since they are often hard to fix if you make a mistake.

4. Output spaces

To force extra spaces into the output, use a tilde “~” for each space you want:

$$x\tilde{=}y\tilde{+}z\tilde{}$$

gives

$$x = y + z$$

You can also use a circumflex “^”, which gives a space half the width of a tilde. It is mainly useful for fine-tuning. Tabs may also be used to position pieces of an expression, but the tab stops must be set by TROFF commands.

5. Symbols, Special Names, Greek

EQN knows some mathematical symbols, some mathematical names, and the Greek alphabet. For example,

$$x=2\pi\text{int}\sin(\omega t)\text{dt}$$

produces

$$x=2\pi\int\sin(\omega t)dt$$

Here the spaces in the input are necessary to tell EQN that *int*, *pi*, *sin* and *omega* are separate entities that should get special treatment. The *sin*, digit 2, and parentheses

are set in roman type instead of italic; *pi* and *omega* are made Greek; and *int* becomes the integral sign.

When in doubt, leave spaces around separate parts of the input. A very common error is to type $f(pi)$ without leaving spaces on both sides of the *pi*. As a result, EQN does not recognize *pi* as a special word, and it appears as $f(pi)$ instead of $f(\pi)$.

A complete list of EQN names appears in section 23. Knowledgeable users can also use TROFF four-character names for anything EQN doesn't know about, like $\backslash bs$ for the Bell System sign $\text{\textcircled{B}}$.

6. Spaces, Again

The only way EQN can deduce that some sequence of letters might be special is if that sequence is separated from the letters on either side of it. This can be done by surrounding a special word by ordinary spaces (or tabs or newlines), as we did in the previous section.

You can also make special words stand out by surrounding them with tildes or circumflexes:

$$x\tilde{=}2\tilde{\pi}\tilde{\text{int}}\tilde{\text{sin}}(\tilde{\omega}\tilde{t})\tilde{\text{dt}}$$

is much the same as the last example, except that the tildes not only separate the magic words like *sin*, *omega*, and so on, but also add extra spaces, one space per tilde:

$$x = 2 \pi \int \sin(\omega t) dt$$

Special words can also be separated by braces { } and double quotes "...", which have special meanings that we will see soon.

7. Subscripts and Superscripts

Subscripts and superscripts are obtained with the words *sub* and *sup*.

$$x\text{sup }2 + y\text{sub }k$$

gives

$$x^2+y_k$$

EQN takes care of all the size changes and vertical motions needed to make the output look right. The words *sub* and *sup* must be surrounded by spaces; $x\text{sub}2$ will give you x_{sub2} instead of x_2 . Furthermore, don't

forget to leave a space (or a tilde, etc.) to mark the end of a subscript or superscript. A common error is to say something like

$$y = (x \text{ sup } 2) + 1$$

which causes

$$y = (x^2) + 1$$

instead of the intended

$$y = (x^2) + 1$$

Subscripted subscripts and superscripted superscripts also work:

$$x \text{ sub } i \text{ sub } 1$$

is

$$x_{i_1}$$

A subscript and superscript on the same thing are printed one above the other if the subscript comes *first*:

$$x \text{ sub } i \text{ sup } 2$$

is

$$x_i^2$$

Other than this special case, *sub* and *sup* group to the right, so *x sup y sub z* means $x^y z$, not x^y_z .

8. Braces for Grouping

Normally, the end of a subscript or superscript is marked simply by a blank (or tab or tilde, etc.) What if the subscript or superscript is something that has to be typed with blanks in it? In that case, you can use the braces { and } to mark the beginning and end of the subscript or superscript:

$$e \text{ sup } \{i \text{ omega } t\}$$

is

$$e^{i\omega t}$$

Rule: Braces can *always* be used to force EQN to treat something as a unit, or just to make your intent perfectly clear. Thus:

$$x \text{ sub } \{i \text{ sub } 1\} \text{ sup } 2$$

is

$$x_{i_1}^2$$

with braces, but

$$x \text{ sub } i \text{ sub } 1 \text{ sup } 2$$

is

$$x_{i_1}^2$$

which is rather different.

Braces can occur within braces if necessary:

$$e \text{ sup } \{i \text{ pi sup } \{\text{rho} + 1\}\}$$

is

$$e^{i\pi^{\rho+1}}$$

The general rule is that anywhere you could use some single thing like *x*, you can use an arbitrarily complicated thing if you enclose it in braces. EQN will look after all the details of positioning it and making it the right size.

In all cases, make sure you have the right number of braces. Leaving one out or adding an extra will cause EQN to complain bitterly.

Occasionally you will have to print braces. To do this, enclose them in double quotes, like "{". Quoting is discussed in more detail in section 14.

9. Fractions

To make a fraction, use the word *over*:

$$a + b \text{ over } 2c = 1$$

gives

$$\frac{a + b}{2c} = 1$$

The line is made the right length and positioned automatically. Braces can be used to make clear what goes over what:

$$\{\alpha + \beta\} \text{ over } \{\sin(x)\}$$

is

$$\frac{\alpha + \beta}{\sin(x)}$$

What happens when there is both an *over* and a *sup* in the same expression? In such an apparently ambiguous case, EQN does the *sup* before the *over*, so

$$-b \text{ sup } 2 \text{ over } \pi$$

is $\frac{-b^2}{\pi}$ instead of $-b^{\frac{2}{\pi}}$. The rules which

decide which operation is done first in cases like this are summarized in section 23. When in doubt, however, use braces to make clear what goes with what.

10. Square Roots

To draw a square root, use *sqrt*:

$\sqrt{a+b} + 1$ over $\sqrt{ax^2 + bx + c}$
is

$$\sqrt{a+b} + \frac{1}{\sqrt{ax^2+bx+c}}$$

Warning — square roots of tall quantities look lousy, because a root-sign big enough to cover the quantity is too dark and heavy:

$$\sqrt{\{a \sup 2 \text{ over } b \sub 2\}}$$

is

$$\sqrt{\frac{a^2}{b_2}}$$

Big square roots are generally better written as something to the power 1/2:

$$(a^2/b_2)^{1/2}$$

which is

$$(a \sup 2 / b \sub 2) \sup \text{half}$$

11. Summation, Integral, Etc.

Summations, integrals, and similar constructions are easy:

sum from $i=0$ to $\{i = \text{inf}\}$ $x \sup i$
produces

$$\sum_{i=0}^{i=\infty} x^i$$

Notice that we used braces to indicate where the upper part $i=\infty$ begins and ends. No braces were necessary for the lower part $i=0$, because it contained no blanks. The braces will never hurt, and if the *from* and *to* parts contain any blanks, you must use braces around them.

The *from* and *to* parts are both optional, but if both are used, they have to occur in that order.

Other useful characters can replace the *sum* in our example:

int prod union inter

become, respectively,

$$\int \prod \cup \cap$$

Since the thing before the *from* can be anything, even something in braces, *from-to* can often be used in unexpected ways:

$$\lim \text{from } \{n \rightarrow \text{inf}\} x \sub n = 0$$

is

$$\lim_{n \rightarrow \infty} x_n = 0$$

12. Size and Font Changes

By default, equations are set in 10-point type (the same size as this guide), with standard mathematical conventions to determine what characters are in roman and what in italic. Although EQN makes a valiant attempt to use esthetically pleasing sizes and fonts, it is not perfect. To change sizes and fonts, use *size n* and *roman*, *italic*, *bold* and *fat*. Like *sub* and *sup*, size and font changes affect only the thing that follows them, and revert to the normal situation at the end of it. Thus

$$\text{bold } x \ y$$

is

$$\mathbf{x} \ \mathbf{y}$$

and

$$\text{size 14 bold } x = y +$$

$$\text{size 14 } \{\alpha + \beta\}$$

gives

$$\mathbf{x} = \mathbf{y} + \alpha + \beta$$

As always, you can use braces if you want to affect something more complicated than a single letter. For example, you can change the size of an entire equation by

$$\text{size 12 } \{ \dots \}$$

Legal sizes which may follow *size* are 6, 7, 8, 9, 10, 11, 12, 14, 16, 18, 20, 22, 24, 28, 36. You can also change the size *by* a given amount; for example, you can say *size +2* to make the size two points bigger.

or *size -3* to make it three points smaller. This has the advantage that you don't have to know what the current size is.

If you are using fonts other than roman, italic and bold, you can say *font X* where *X* is a one character TROFF name or number for the font. Since EQN is tuned for roman, italic and bold, other fonts may not give quite as good an appearance.

The *fat* operation takes the current font and widens it by overstriking: *fat grad* is ∇ and *fat {x sub i}* is x_i .

If an entire document is to be in a non-standard size or font, it is a severe nuisance to have to write out a size and font change for each equation. Accordingly, you can set a "global" size or font which thereafter affects all equations. At the beginning of any equation, you might say, for instance,

```
.EQ
  gsize 16
  gfont R
...
.EN
```

to set the size to 16 and the font to roman thereafter. In place of R, you can use any of the TROFF font names. The size after *gsize* can be a relative change with + or -.

Generally, *gsize* and *gfont* will appear at the beginning of a document but they can also appear throughout a document: the global font and size can be changed as often as needed. For example, in a footnote† you will typically want the size of equations to match the size of the footnote text, which is two points smaller than the main text. Don't forget to reset the global size at the end of the footnote.

† Like this one, in which we have a few random expressions like χ_i and π^2 . The sizes for these were set by the command *gsize -2*.

13. Diacritical Marks

To get funny marks on top of letters, there are several words:

x dot	\dot{x}
x dotdot	\ddot{x}
x hat	\hat{x}
x tilde	\tilde{x}
x vec	\vec{x}
x dyad	\overline{x}
x bar	\bar{x}
x under	\underline{x}

The diacritical mark is placed at the right height. The *bar* and *under* are made the right length for the entire construct, as in $\overline{x+y+z}$; other marks are centered.

14. Quoted Text

Any input entirely within quotes ("...") is not subject to any of the font changes and spacing adjustments normally done by the equation setter. This provides a way to do your own spacing and adjusting if needed:

italic "sin(x)" + sin (x)

is

sin(x)+sin(x)

Quotes are also used to get braces and other EQN keywords printed:

"{ size alpha }"

is

{ *size alpha* }

and

roman "{ size alpha }"

is

{ size alpha }

The construction "" is often used as a place-holder when grammatically EQN needs something, but you don't actually want anything in your output. For example, to make ²He, you can't just type *sup 2 roman He* because a *sup* has to be a superscript *on*

something. Thus you must say

```
" sup 2 roman He
```

To get a literal quote use `"\"`. TROFF characters like `\(b`s can appear unquoted, but more complicated things like horizontal and vertical motions with `\h` and `\v` should always be quoted. (If you've never heard of `\h` and `\v`, ignore this section.)

15. Lining Up Equations

Sometimes it's necessary to line up a series of equations at some horizontal position, often at an equals sign. This is done with two operations called *mark* and *lineup*.

The word *mark* may appear once at any place in an equation. It remembers the horizontal position where it appeared. Successive equations can contain one occurrence of the word *lineup*. The place where *lineup* appears is made to line up with the place marked by the previous *mark* if at all possible. Thus, for example, you can say

```
.EQ I
x+y mark = z
.EN
.EQ I
x lineup = 1
.EN
```

to produce

```
x+y=z
x=1
```

For reasons too complicated to talk about, when you use EQN and '-ms', use either EQ1 or EQ1. *mark* and *lineup* don't work with centered equations. Also bear in mind that *mark* doesn't look ahead;

```
x mark = 1
...
x+y lineup = z
```

isn't going to work, because there isn't room for the `x+y` part after the *mark* remembers where the `x` is.

16. Big Brackets, Etc.

To get big brackets `[]`, braces `{ }`, parentheses `()`, and bars `||` around things, use the *left* and *right* commands:

```
left { a over b + 1 right }
~ = ~ left ( c over d right )
+ left [ e right ]
```

is

$$\left\{ \frac{a}{b} + 1 \right\} = \left(\frac{c}{d} \right) + [e]$$

The resulting brackets are made big enough to cover whatever they enclose. Other characters can be used besides these, but they are not likely to look very good. One exception is the *floor* and *ceiling* characters:

```
left floor x over y right floor
< = left ceiling a over b right ceiling
```

produces

$$\left\lfloor \frac{x}{y} \right\rfloor \leq \left\lceil \frac{a}{b} \right\rceil$$

Several warnings about brackets are in order. First, braces are typically bigger than brackets and parentheses, because they are made up of three, five, seven, etc., pieces, while brackets can be made up of two, three, etc. Second, big left and right parentheses often look poor, because the character set is poorly designed.

The *right* part may be omitted: a "left something" need not have a corresponding "right something". If the *right* part is omitted, put braces around the thing you want the left bracket to encompass. Otherwise, the resulting brackets may be too large.

If you want to omit the *left* part, things are more complicated, because technically you can't have a *right* without a corresponding *left*. Instead you have to say

```
left "" ..... right )
```

for example. The *left* "" means a "left nothing". This satisfies the rules without hurting your output.

17. Piles

There is a general facility for making vertical piles of things; it comes in several flavors. For example:

```
A ~ = ~ left [
  pile { a above b above c }
  ~ ~ pile { x above y above z }
right ]
```

will make

$$A = \begin{bmatrix} a & x \\ b & y \\ c & z \end{bmatrix}$$

The elements of the pile (there can be as many as you want) are centered one above another, at the right height for most purposes. The keyword *above* is used to separate the pieces; braces are used around the entire list. The elements of a pile can be as complicated as needed, even containing more piles.

Three other forms of pile exist: *lpile* makes a pile with the elements left-justified; *rpile* makes a right-justified pile; and *cpile* makes a centered pile, just like *pile*. The vertical spacing between the pieces is somewhat larger for *l-*, *r-* and *cpiles* than it is for ordinary piles.

```
roman sign (x) ~ = ~
left {
  lpile { 1 above 0 above -1 }
  ~ ~ lpile
  { if x > 0 above if x = 0 above if x < 0 }
```

makes

$$\text{sign}(x) = \begin{cases} 1 & \text{if } x > 0 \\ 0 & \text{if } x = 0 \\ -1 & \text{if } x < 0 \end{cases}$$

Notice the left brace without a matching right one.

18. Matrices

It is also possible to make matrices. For example, to make a neat array like

$$\begin{matrix} x_i & x^2 \\ y_i & y^2 \end{matrix}$$

you have to type

```
matrix {
  ccol { x sub i above y sub i }
  ccol { x sup 2 above y sup 2 }
}
```

This produces a matrix with two centered columns. The elements of the columns are then listed just as for a pile, each element separated by the word *above*. You can also use *lcol* or *rcol* to left or right adjust columns. Each column can be separately adjusted, and there can be as many columns as you like.

The reason for using a matrix instead of two adjacent piles, by the way, is that if the elements of the piles don't all have the same height, they won't line up properly. A matrix forces them to line up, because it looks at the entire structure before deciding what spacing to use.

A word of warning about matrices — *each column must have the same number of elements in it*. The world will end if you get this wrong.

19. Shorthand for In-line Equations

In a mathematical document, it is necessary to follow mathematical conventions not just in display equations, but also in the body of the text, for example by making variable names like *x* italic. Although this could be done by surrounding the appropriate parts with *.EQ* and *.EN*, the continual repetition of *.EQ* and *.EN* is a nuisance. Furthermore, with '-ms', *.EQ* and *.EN* imply a displayed equation.

.EQN provides a shorthand for short in-line expressions. You can define two characters to mark the left and right ends of an in-line equation, and then type expressions right in the middle of text lines. To set both the left and right characters to dollar signs, for example, add to the beginning of your document the three lines

```
.EQ
delim $$
.EN
```

Having done this, you can then say things like

Let α_i be the primary variable, and let β be zero. Then we can show that x_i is ≥ 0 .

This works as you might expect — spaces, newlines, and so on are significant in the text, but not in the equation part itself. Multiple equations can occur in a single input line.

Enough room is left before and after a line that contains in-line expressions that something like $\sum_{i=1}^n x_i$, does not interfere with the lines surrounding it.

To turn off the delimiters,

```
.EQ
delim off
.EN
```

Warning: don't use braces, tildes, circumflexes, or double quotes as delimiters — chaos will result.

20. Definitions

EQN provides a facility so you can give a frequently-used string of characters a name, and thereafter just type the name instead of the whole string. For example, if the sequence

$$x_{i-1} + y_{i-1}$$

appears repeatedly throughout a paper, you can save re-typing it each time by defining it like this:

```
define xy `x sub i sub 1 + y sub i sub 1`
```

This makes *xy* a shorthand for whatever characters occur between the single quotes in the definition. You can use any character instead of quote to mark the ends of the definition, so long as it doesn't appear inside the definition.

Now you can use *xy* like this:

```
.EQ
f(x) = xy ...
.EN
```

and so on. Each occurrence of *xy* will expand into what it was defined as. Be careful to leave spaces or their equivalent

around the name when you actually use it, so EQN will be able to identify it as special.

There are several things to watch out for. First, although definitions can use previous definitions, as in

```
.EQ
define xi `x sub i`
define xil `xi sub 1`
.EN
```

don't define something in terms of itself" A favorite error is to say

```
define X `roman X`
```

This is a guaranteed disaster, since *X* is now defined in terms of itself. If you say

```
define X `roman "X" `
```

however, the quotes protect the second *X*, and everything works fine.

EQN keywords can be redefined. You can make / mean *over* by saying

```
define / `over`
```

or redefine *over* as / with

```
define over `/`
```

If you need different things to print on a terminal and on the typesetter, it is sometimes worth defining a symbol differently in NEQN and EQN. This can be done with *ndefine* and *tdefine*. A definition made with *ndefine* only takes effect if you are running NEQN; if you use *tdefine*, the definition only applies for EQN. Names defined with plain *define* apply to both EQN and NEQN.

21. Local Motions

Although EQN tries to get most things at the right place on the paper, it isn't perfect, and occasionally you will need to tune the output to make it just right. Small extra horizontal spaces can be obtained with tilde and circumflex. You can also say *back n* and *fwd n* to move small amounts horizontally. *n* is how far to move in 1/100's of an em (an em is about the width of the letter 'm'.) Thus *back 50* moves back about half the width of an m. Similarly you can move things up or down with *up n* and *down n*. As with *sub* or *sup*, the local motions affect the

next thing in the input, and this can be something arbitrarily complicated if it is enclosed in braces.

As an example of local motions, consider tucking the limits in under an integral sign. Normally if you say

```
int sub 0 sup 1
```

it looks like

$$\int_0^1$$

which is awful. The intuitively appealing

```
int from 0 to 1
```

is

$$\int_0^1$$

which is not normally used. But if you say

```
int sub back 40 down 50 0 sup up 30 1
```

you get

$$\int_0^1$$

(These values are determined experimentally.) Of course this is a nuisance to type, so you would first make definitions like this:

```
tdefine lower `sub back 40 down 50 `
tdefine upper `sup up 30 `
```

and then say

```
int lower 0 upper 1
```

22. A Large Example

Here is the complete source for the three display equations in the abstract of this guide.

```
EQ 1
G(z) mark = e sup { ln G(z) }
= exp left (
sum from k >= 1 { S sub k z sup k } over k right )
= prod from k >= 1 e sup { S sub k z sup k / k }
EN
EQ 1
lineup = left ( 1 + S sub 1 z +
{ S sub 1 sup 2 z sup 2 } over 2! + ... right )
left ( 1 + { S sub 2 z sup 2 } over 2
+ { S sub 2 sup 2 z sup 4 } over { 2 sup 2 cdot 2! }
+ ... right )
EN
EQ 1
lineup = sum from m >= 0 left (
```

```
sum from
pile { k sub 1 , k sub 2 , ... , k sub m } >= 0
above
k sub 1 + 2k sub 2 + ... + mk sub m = m }
{ S sub 1 sup { k sub 1 } } over { 1 sup k sub 1 k sub 1 ! } }
{ S sub 2 sup { k sub 2 } } over { 2 sup k sub 2 k sub 2 ! } }
...
{ S sub m sup { k sub m } } over { m sup k sub m k sub m ! }
right ) z sup m
EN
```

23. Keywords, Precedences, Etc.

If you don't use braces, EQN will do operations in the order shown in this list.

```
dyad vec under bar tilde hat dot dotdot
fwd back down up
fat roman italic bold size
sub sup sqrt over
from to
```

These operations group to the left:

```
over sqrt left right
```

All others group to the right.

Digits, parentheses, brackets, punctuation marks, and these mathematical words are converted to Roman font when encountered:

```
sin cos tan sinh cosh tanh arc
max min lim log ln exp
Re Im and if for det
```

These character sequences are recognized and translated as shown.

```
>= >=
<= <=
== ==
!= !=
+- +-
-> ->
<- <-
<< <<
>> >>
inf inf
partial partial
half half
prime prime
approx approx
nothing nothing
cdot cdot
times times
del del
```

grad	∇
...	...
sum	Σ
int	\int
prod	Π
union	\cup
inter	\cap

fwd	21	to	11
gfont	12	under	13
gsize	12	up	21
hat	13	vec	13
italic	12	\sim , $\hat{}$	4, 6
lcol	18	{ }	8
left	16	"..."	8, 14
lineup	15		

To obtain Greek letters, simply spell them out in whatever case you want:

DELTA	Δ	iota	ι
GAMMA	Γ	kappa	κ
LAMBDA	Λ	lambda	λ
OMEGA	Ω	mu	μ
PHI	Φ	nu	ν
PI	Π	omega	ω
PSI	Ψ	omicron	\omicron
SIGMA	Σ	phi	ϕ
THETA	Θ	pi	π
UPSILON	Υ	psi	ψ
XI	Ξ	rho	ρ
alpha	α	sigma	σ
beta	β	tau	τ
chi	χ	theta	θ
delta	δ	upsilon	υ
epsilon	ϵ	xi	ξ
eta	η	zeta	ζ
gamma	γ		

These are all the words known to EQN (except for characters with names), together with the section where they are discussed.

above	17, 18	lpile	17
back	21	mark	15
bar	13	matrix	18
bold	12	ndefine	20
ccol	18	over	9
col	18	pile	17
cpile	17	rcol	18
define	20	right	16
delim	19	roman	12
dot	13	rpile	17
dotdot	13	size	12
down	21	sqrt	10
dyad	13	sub	7
fat	12	sup	7
font	12	tdefine	20
from	11	tilde	13

24. Troubleshooting

If you make a mistake in an equation, like leaving out a brace (very common) or having one too many (very common) or having a *sup* with nothing before it (common), EQN will tell you with the message

syntax error between lines x and y, file z

where *x* and *y* are approximately the lines between which the trouble occurred, and *z* is the name of the file in question. The line numbers are approximate — look nearby as well. There are also self-explanatory messages that arise if you leave out a quote or try to run EQN on a non-existent file.

If you want to check a document before actually printing it (on UNIX only),

```
eqn files >/dev/null
```

will throw away the output but print the messages.

If you use something like dollar signs as delimiters, it is easy to leave one out. This causes very strange troubles. The program *checkeq* (on GCOS, use *.checkeq* instead) checks for misplaced or missing dollar signs and similar troubles.

In-line equations can only be so big because of an internal buffer in TROFF. If you get a message "word overflow", you have exceeded this limit. If you print the equation as a displayed equation this message will usually go away. The message "line overflow" indicates you have exceeded an even bigger buffer. The only cure for this is to break the equation into two separate ones.

On a related topic, EQN does not break equations by itself — you must split long equations up across multiple lines by yourself, marking each by a separate `.EQEN`

sequence. EQN does warn about equations that are too long to fit on one line.

25. Use on UNIX

To print a document that contains mathematics on the UNIX typesetter,

```
eqn files | troff
```

If there are any TROFF options, they go after the TROFF part of the command. For example,

```
eqn files | troff -ms
```

To run the same document on the GCOS typesetter, use

```
eqn files | troff -g (other options) | gcat
```

A compatible version of EQN can be used on devices like teletypes and DASI and GSI terminals which have half-line forward and reverse capabilities. To print equations on a Model 37 teletype, for example, use

```
neqn files | nroff
```

The language for equations recognized by NEQN is identical to that of EQN, although of course the output is more restricted.

To use a GSI or DASI terminal as the output device,

```
neqn files | nroff -Tx
```

where x is the terminal type you are using, such as 300 or 300S.

EQN and NEQN can be used with the TBL program[2] for setting tables that contain mathematics. Use TBL before [N]EQN, like this:

```
tbl files | eqn | troff
tbl files | neqn | nroff
```

26. Use on GCOS

This space intentionally left blank


27. Acknowledgments


We are deeply indebted to J. F. Ossanna, the author of TROFF, for his willingness to extend TROFF to make our task easier, and for his continuous assistance during the development and evolution of EQN. We are also grateful to A. V. Aho for advice on language design, to S. C. Johnson for assistance with the YACC compiler-compiler, and to all the EQN users who have made helpful suggestions and criticisms.

References

- [1] J. F. Ossanna, "TROFF User's Manual", Bell Laboratories internal memorandum.
- [2] M. E. Lesk, "Typing Documents on UNIX", Bell Laboratories internal memorandum.
- [3] M. E. Lesk, "TBL— A Program for Setting Tables", Bell Laboratories internal memorandum.







New Graphic Symbols for EQN and NEQN

Carmela Scrocca

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

There is now available on UNIX and GCOS a set of special characters frequently used in technical typing. In the past, authors have sometimes written out these symbols in English; others just assumed their secretary or typist had these symbols ready and waiting. These characters, however, are not part of the standard terminal or typesetter character sets, but are built-up of those already available. They can presently be produced for phototypesetter output by using EQN/TROFF; NEQN/NROFF can be used for computer terminal output.

This document displays these characters, shows how to use them, and discusses what is involved in making a special character.



New Graphic Symbols for EQN and NEQN

Carmela Scrocca

Bell Laboratories
Murray Hill, New Jersey 07974

Introduction

There is now available on UNIX and GCOS a set of special characters frequently used in technical typing. These characters supplement the ones that come with the typesetter and terminal which both have their own set of standard characters. These special characters are accessed through the math typesetting programs NEQN and EQN.¹ Processed through the NROFF and TROFF² formatting programs, the characters can be output on either a computer terminal or a phototypesetter. Using various NROFF/TROFF conventions, two or more of these existing characters can be built-up and pieced together to draw a new character.

Sections 1 and 2 of this document give a list of these characters and tell how to access and use them. In Sections 3 and 4, the reader will see what is involved in making a special character for both phototypesetter and computer terminal output.

1. The Characters

Table 1 gives a list of the characters, their meanings, and the names by which EQN recognizes them.

The user should be aware that these special characters are not built into the NEQN/EQN program, but are stored in directory /usr/pub; the filename is *eqnchar*. In order to use any of these symbols this file will have to be referenced. This can be done by using file *eqnchar* as the first filename in your output command, such as

```
neqn /usr/pub/eqnchar filenames | nroff
```

for computer terminal output. For phototypesetter output use EQN/TROFF instead of NEQN/NROFF. On GCOS, the characters are in file *./eqnchar*.

Some users will find a constant need for only a few of the special characters. It may be convenient for these users to take their few selected characters, and write them into a file in their own directory. Other users may only need a few characters for use in one particular document. They could edit /usr/pub/eqnchar, copy the desired character definitions into a separate file, and read that file into the beginning of their document file.

Appendices 1a and 1b provide additional characters with their corresponding EQN names. The characters list in Appendix 1a are from the phototypesetter character set and have been assigned EQN names. /usr/pub/eqnchar must be referenced in the output command to use characters in this set. Appendix 1b contains a list of characters already built into EQN and can be used directly with the program.

character name	character	EQN name
sum of two elements	\oplus	ciplus
product of two elements	\otimes	citimes
is congruent to	\cong	=wig
approximately equal to	\doteq	=dot
equals by definition	\triangleq	=del
large star	\bigstar	bigstar
centered star	\star	star
or	\vee	orsign
and	\wedge	andsign
for all	\forall	oppA
there exists	\exists	oppE
is included in	\sqsubset	incl
not a member of	\notin	nomem
angstrom	\AA	angstrom
less than or approximately equal to	\lesseqgtr	<wig
greater than or approximately equal to	\gtrless	>wig
not less than	\nlessgtr	<
not greater than	\ngtrless	>
left angle bracket	\langle	langle
right angle bracket	\rangle	rangle
hbar	\hbar	hbar
parallel	\parallel	
perpendicular	\perp	ppd
angle	\sphericalangle	ang
right angle	\rightangle	rang
implies and is implied by	\longleftrightarrow	<->
implies and is implied by	\Leftrightarrow	<=>
vertical ellipsis	\vdots	3dot
therefore	\therefore	thf

Table 1. The Special Characters

2. Usage

The reader should be familiar with EQN as these characters work with the program and are used in the same manner as any other mathematical symbol. An EQN name must be separated from surrounding input (with spaces) in order to be recognized as a special character. For example, just as you would say

$$\pi = \sum x^i$$

to get

$$\pi = \sum x'$$

you could say

$$x \text{ sub } 3 \text{ } \overset{\sim}{=} \text{ wig } \pi \text{ star } y \text{ sub } \text{ ppd}$$

to give you

$$x_3 \cong \pi * y_{\perp}$$

of which *wig*, *star* and *ppd* are a few of the new special characters. These symbols will work in both displayed and in-line equations. There are no bold or italic versions of the symbols.

3. Creating Special Characters for EQN

The symbols discussed here were created by taking pre-existing characters and piecing them together with the input and output conventions and escape sequences made available through NROFF and TROFF. Most commonly used here were the local horizontal and vertical motions, overstrike and zero width functions, and the point-size change function. These are used primarily for phototypesetter output; making characters for computer terminal output will be discussed in Section 4. The definitions of all the special characters are listed in Appendix 2.

Local horizontal and vertical motions are used to move a character up, down, or to the left or right depending on where you want to position your character. These motions are generated by the escape sequences `\u`, `\d`, `\r`, `\v` and `\h`. The motions are expressed in terms of ems; an em is approximately the width of the letter 'm'. By using ems, the amount of motion will always be in proportion to character size. The `\u` and `\d` sequences give vertical motions of 1/2 em up and down, respectively; `\r` gives an upward motion of 1 em. The `\v` (vertical motion) and `\h` (horizontal motion) escape sequences allow you to move any fraction of an em. The distance must be enclosed in ' marks and the direction of movement can be indicated by making it either positive or negative. For example, if you wanted a downward vertical motion of 3/10 of an em, you would say `\v'.3m'`, and an upward vertical motion of 6/10 of an em would be `\v'-.6m'`. The same basic rules apply to horizontal motions where positive moves to the right, and negative to the left. The "is much greater than" symbol `>>` shows a simple horizontal motion:

$$>\backslash h'-.3m'>$$

(`>>` is not a special character, but built into EQN; see Appendix 1b.)

The overstrike function `\o` simply overprints characters on top of another, centered on the widest character. This function was used to create the "sum of two elements" symbol \oplus by saying

$$\backslash o^{(p)\backslash (c) i'$$

where `\(p` and `\(c` are the escapes for + and O. The string of characters to be overstruck must be enclosed in ' marks. When using the overstrike function be sure not to use any motions or it will not work.

Similar to the overstrike function is the bracket building function `\b`. Instead of centering one character on the other, the bracket building function piles the characters vertically. The "angle brackets" were built-up using the `\b` function:

$$\begin{array}{l} \backslash s-3\backslash b\backslash (s\backslash e'\backslash s0 \\ \backslash s-3\backslash b\backslash e\backslash (s\backslash s0 \end{array} \left. \begin{array}{l} \langle \\ \rangle \end{array} \right.$$

(\e and \s) are the escapes for \ and /.) The first character in the string is positioned at the top of the pile, and on down with the last character at the bottom. \s is the escape sequence for a size change; we'll get to this in a while.

The zero width function \z enables you to print a character without moving after it is printed. \z often makes it easier to position your next character rather than figuring out where the first one moved you to. This function is used with \zn where n is the character to be printed. \z can only be applied to one character at a time so the sequence would have to be repeated. An example of this can be the definition of the "vertical ellipsis" ; where \z forces the dots to remain in place; otherwise, horizontal motions would be needed to realign the dots.

$$\backslash v'-.8m\backslash z.\backslash v'.5m\backslash z.\backslash v'.5m'.\backslash v'-.2m'$$

The zero width function can also be used to darken a character: \zxx will print an x, stay in place, and print it again.

The point-size change function is used frequently to make one character fit in proportion to another. The escape sequence for this function is \s. The change in size would be indicated by however many point sizes you want to change. \s-2 will cause a reduction of 2 point sizes while \s+2 will enlarge by 2 point sizes. For example, in making the "equals by definition" symbol $\stackrel{\Delta}{=}$, the Δ was made slightly smaller to fit comfortably over the =. This was done by

$$\backslash v'.3m\backslash z=\backslash v'-.6m\backslash h'.3m\backslash s-1\backslash (*D\backslash s+1\backslash v'.3m'$$

hence reducing the Δ (\(*D is the escape for this) by 1 point size before printing it, and then returning back to the previous size. \s0 can also be used to bring you back to the previous size.

Another handy tool is the font change function \f. Since EQN automatically sets its characters in italics, it will be necessary to specify any other font. To change fonts, use \fx where x is the desired font (R for roman, B for bold, I for italic, etc.). \fP will revert back to the previous font. A prime example of this would be the "angstrom" symbol \AA :

$$\backslash fR\backslash zA\backslash v'-.3m\backslash h'.2m\backslash (de\backslash v'.3m\backslash fP\backslash h'.2m'$$

where the capital "A" is always printed in the Roman font.

A few words of caution. If in building your character you have used vertical motions or point-size or font changes, you must remember to undo them, or whatever follows will be off the main line or in the wrong size or font. It may also be necessary at times to use a horizontal motion to ensure enough space before and after your character. Also, you should test your character before putting it to actual use. Try it out with changes in original font or point size and see how it reacts to the change; don't be surprised if your character falls apart.

When you are satisfied with your character and it is ready to use, introduce it into your file with the define facility provided by EQN. For example, you would define the "not greater than" symbol (∇) as

$$\text{define } \nabla \text{ \% "\o'>\(or'" \%}$$

Now by using ∇ in an equation, you will get ∇ .

In the previous example, the %'s and "'s are two different kinds of delimiters. The outside pair are essential to mark the beginning and end of the definition. (%'s were used, although any character will do as long as it is not used in the definition.) Also, any definition containing TROFF commands should be enclosed in " marks, so EQN will treat the TROFF commands as a unit.

4. Creating Special Characters for NEQN

To get a special character to print out on a computer terminal is not quite as involved as on the phototypesetter. You are restricted to working with only the characters available on the print wheel and movement is limited. Vertical motion can be obtained by `\u` and `\d` but this will only give you a motion of $\frac{1}{2}$ line space per escape sequence. Spacing and backspacing is about all the horizontal motion you'll get. You can, however, use NEQN to define a character as was done with the "less than or approximately equal to" symbol \lesssim :

```
ndefine <wig % < from "" %
```

All TROFF conventions work in NROFF, but because of the lack of characters, they may produce strange effects when using NEQN/NROFF. Therefore, it may be necessary to separately define characters for phototypesetter and computer terminal output. *tdefine* applies to definitions for EQN; *ndefine* works with NEQN. *define* uses the same definition for both.

Built-up characters for terminal output are usually just good enough for identification's sake; generally, they look lousy. However, given the time and additional effort, these characters can be refined so that their output on the terminal is quite satisfactory.

Acknowledgements

I would like to thank J. F. Ossanna, M. E. Lesk and other members of Center 127 for all their help; and special thanks to B. W. Kernighan for his teaching, guidance and patience throughout. I am especially grateful to S. P. Morgan for his encouragement and shared enthusiasm; without his aid and the concurrence of Dept. 7133K supervision, none of this would have been possible.

References

1. "A System for Typesetting Mathematics," B. W. Kernighan and L. L. Cherry, Computing Science Technical Report #17.
2. NROFF/TROFF User's Manual, J. F. Ossanna, BTL internal memorandum.

Appendix 1a

Additional Symbols from Phototypesetter Character Set

character	EQN name
$\frac{1}{4}$	quarter
$\frac{3}{4}$	3quarter
$^{\circ}$	degree
\square	square
\bigcirc	circle
\blacksquare	blot
\bullet	bullet
\approx	-wig
\sim	wig
$\&$	prop
\emptyset	empty
\in	member
\cup	cup
* \supset	cap
* \subset	subset
* \supseteq	supset
* \supsetneq	!subset
* \supsetneq	!supset

* NEQN/NROFF does not produce these symbols.

Appendix 1b

Additional Symbols Provided by EQN

character	EQN name
* $\frac{1}{2}$	half
\approx	approx
\gg	$> =$
\ll	$< =$
\ggg	$>>$
\lll	$<<$
\rightarrow	$->$
\leftarrow	$<-$
\equiv	$= =$
\neq	$! =$
\pm	$+ -$
∞	inf
∂	partial
\prime	prime
\cdot	cdot
\times	times
∇	grad
\dots	...
\sum	sum
\int	int
\prod	prod
\cup	union
* \cap	inter

* NEQN/NROFF does not produce these symbols.


Note: — represents a backspace in undefined characters.


EQ

```

tdefine ciplus % "\o\{pl\{ci" %
ndefine ciplus % O--+ %
tdefine citimes % "\o\{mu\{ci" %
ndefine citimes % O-x %
tdefine =wig % "\{eq\h'-\w\{eq'u-\w\{s-2\{ap'u/2u\{v'-.4m\{s-2\{z\{ap\{ap\{s+2\{v'.4m\{h\{w\{eq'u-\w\{s-2\{ap'u/2u" %
ndefine =wig % =-"" %
tdefine bigstar % "\o\{pl\{mu" %
ndefine bigstar % X-|--%
tdefine =dot % "\z\{eq\{v'-.6m\{h'.2m\{s+2\{s-2\{v'.6m\{h'.1m" %
ndefine =dot % = dot %
tdefine orsign % "\s-2\{v'-.15m\{z\{e\{e\{h'-.05m\{z\{s\{s\{h'-.05m\{z\{e\{e\{v'.15m\{s+2" %
ndefine orsign % \e/ %
tdefine andsign % "\s-2\{v'-.15m\{z\{s\{s\{h'-.05m\{z\{e\{e\{v'.15m\{s+2" %
ndefine andsign % \e %
tdefine =del % "\v'.3m\{z=\{v'-.6m\{h'.3m\{s-1\{D\{s+1\{v'.3m" %
ndefine =del % = to DELTA %
tdefine oppA % "\s-2\{v'-.15m\{z\{e\{e\{h'-.05m\{z\{s\{s\{v'-.15m\{h'-.75m\{z-\{z-\{h'.2m\{z-\{z-\{v'.3m\{h'.4m\{s+2" %
ndefine oppA % V---%
tdefine oppE % "\s-3\{v'.2m\{z\{em\{v'-.5m\{z\{em\{v'-.5m\{z\{em\{v'.55m\{h'.9m\{z\{br\{z\{br\{v'.25m\{s+3" %
ndefine oppE % E--/ %
tdefine incl % "\s-1\{z\{or\{h'-.1m\{v'-.45m\{z\{em\{v'.7m\{z\{em\{v'.2m\{em\{v'-.45m\{s+1" %
ndefine incl % C-- %
tdefine nomem % "\o\{mo\{s\{l" %
ndefine nomem % C-- --/ %
tdefine angstrom % "\{R\{zA\{v'-.3m\{h'.2m\{de\{v'.3m\{fP\{h'.2m" %
ndefine angstrom % A to o %
tdefine star % { roman "\v'.5m\{s+3*\{s-3\{v'-.5m"} %
ndefine star % * %
tdefine || % \or\or %
tdefine <wig % "\z<\v'.4m\{ap\{v'-.4m" %
ndefine <wig % { < from "" } %
tdefine >wig % "\z>\v'.4m\{ap\{v'-.4m" %
ndefine >wig % { > from "" } %
tdefine langle % "\s-3\{b\{s\{e\{s0" %
ndefine langle % <%
tdefine rangle % "\s-3\{b\{e\{s\{s0" %
ndefine rangle % >%
tdefine hbar % "\zh\{v'-.6m\{h'.05m\{ru\{v'.6m" %
ndefine hbar % h-\u-\d %
ndefine ppd % _-| %
tdefine ppd % "\o\{ru\{s-2\{or\{s+2" %
tdefine <-> % "\o\{<-\{>->" %
ndefine <-> % "<--->" %
tdefine <=> % "\s-2\{z<\v'.05m\{h'.2m\{z=\{h'.55m\{h'-.6m\{v'-.05m\{s+2" %
ndefine <=> % "<=>" %
tdefine |< % "\o\{<\{or" %
ndefine |< % <-| %
tdefine |> % "\o\{>\{or" %
ndefine |> % |-> %
tdefine ang % "\v'-.15m\{z\{s-2\{s\{s+2\{v'.15m\{ru" %
ndefine ang % /-_%
tdefine rang % "\z\{or\{h'.15m\{ru" %
ndefine rang % L %
tdefine 3dot % "\v'-.8m\{z.\v'.5m\{z.\v'.5m\{v'-.2m" %
ndefine 3dot % .-\u.-\u.\d\ %
tdefine thf % "\v'-.5m\{v'.5m" %
ndefine thf % ..-\u.\d %
EN

```






The PWB/UNIX* document entitled:
PWBIUNIX View Graph and Slide Macros
is not yet available.

* UNIX is a Trademark/Service Mark of the Bell System.







BC — An Arbitrary Precision Desk-Calculator Language

Lorinda Cherry

Robert Morris

Bell Laboratories,
Murray Hill, New Jersey 07974

ABSTRACT

BC is a language and a compiler for doing arbitrary precision arithmetic on the PDP-11 under the UNIX time-sharing system. The output of the compiler is interpreted and executed by a collection of routines which can input, output, and do arithmetic on indefinitely large integers and on scaled fixed-point numbers.

These routines are themselves based on a dynamic storage allocator. Overflow does not occur until all available core storage is exhausted.

The language has a complete control structure as well as immediate-mode operation. Functions can be defined and saved for later execution.

Two five hundred-digit numbers can be multiplied to give a thousand digit result in about ten seconds.

A small collection of library functions is also available, including sin, cos, arctan, log, exponential, and Bessel functions of integer order.

Some of the uses of this compiler are

- to do computation with large integers,
- to do computation accurate to many decimal places,
- conversion of numbers from one base to another base.



BC -- An Arbitrary Precision Desk-Calculator Language

Lorinda Cherry

Robert Morris

Bell Laboratories,
Murray Hill, New Jersey 07974

Introduction

BC is a language and a compiler for doing arbitrary precision arithmetic on the UNIX time-sharing system [1]. The compiler was written to make conveniently available a collection of routines (called DC [6]) which are capable of doing arithmetic on integers of arbitrary size. The compiler is by no means intended to provide a complete programming language. It is a minimal language facility.

There is a scaling provision that permits the use of decimal point notation. Provision is made for input and output in bases other than decimal. Numbers can be converted from decimal to octal by simply setting the output base to equal 8.

The actual limit on the number of digits that can be handled depends on the amount of storage available on the machine. Manipulation of numbers with many hundreds of digits is possible even on the smallest versions of UNIX.

The syntax of BC has been deliberately selected to agree substantially with the C language [2,3]. Those who are familiar with C will find few surprises in this language.

Simple Computations with Integers

The simplest kind of statement is an arithmetic expression on a line by itself. For instance, if you type in the line:

```
142857 + 285714
```

the program responds immediately with the line

```
428571
```

The operators $-$, $*$, $/$, $\%$, and $^$ can also be used; they indicate subtraction, multiplication, division, remaindering, and exponentiation, respectively. Division of integers produces an integer result truncated toward zero. Division by zero produces an error comment.

Any term in an expression may be prefixed by a minus sign to indicate that it is to be negated (the 'unary' minus sign). The expression

```
7+-3
```

is interpreted to mean that -3 is to be added to 7.

More complex expressions with several operators and with parentheses are interpreted just as in Fortran, with $^$ having the greatest binding power, then $*$ and $\%$ and $/$, and finally $+$ and $-$. Contents of parentheses are evaluated before material outside the parentheses. Exponentiations are performed from right to left and the other operators from left to right. The two expressions

$a^b \cdot c$ and $a^{(b \cdot c)}$

are equivalent, as are the two expressions

$a \cdot b \cdot c$ and $(a \cdot b) \cdot c$

BC shares with Fortran and C the undesirable convention that

$a/b \cdot c$ is equivalent to $(a/b) \cdot c$

Internal storage registers to hold numbers have single lower-case letter names. The value of an expression can be assigned to a register in the usual way. The statement

```
x = x + 3
```

has the effect of increasing by three the value of the contents of the register named x. When, as in this case, the outermost operator is an =, the assignment is performed but the result is not printed. Only 26 of these named storage registers are available.

There is a built-in square root function whose result is truncated to an integer (but see scaling below). The lines

```
x = sqrt(191)
```

```
x
```

produce the printed result

```
13
```

Bases

There are special internal quantities, called 'ibase' and 'obase'. The contents of 'ibase', initially set to 10, determines the base used for interpreting numbers read in. For example, the lines

```
ibase = 8
```

```
11
```

will produce the output line

```
9
```

and you are all set up to do octal to decimal conversions. Beware, however of trying to change the input base back to decimal by typing

```
ibase = 10
```

Because the number 10 is interpreted as octal, this statement will have no effect. For those who deal in hexadecimal notation, the characters A-F are permitted in numbers (no matter what base is in effect) and are interpreted as digits having values 10-15 respectively. The statement

```
ibase = A
```

will change you back to decimal input base no matter what the current input base is. Negative and large positive input bases are permitted but useless. No mechanism has been provided for the input of arbitrary numbers in bases less than 1 and greater than 16.

The contents of 'obase', initially set to 10, are used as the base for output numbers. The lines

```
obase = 16
1000
```

will produce the output line

```
3E8
```

which is to be interpreted as a 3-digit hexadecimal number. Very large output bases are permitted, and they are sometimes useful. For example, large numbers can be output in groups of five digits by setting 'obase' to 100000. Strange (i.e. 1, 0, or negative) output bases are handled appropriately.

Very large numbers are split across lines with 70 characters per line. Lines which are continued end with \. Decimal output conversion is practically instantaneous, but output of very large numbers (i.e., more than 100 digits) with other bases is rather slow. Non-decimal output conversion of a one hundred digit number takes about three seconds.

It is best to remember that 'ibase' and 'obase' have no effect whatever on the course of internal computation or on the evaluation of expressions, but only affect input and output conversion, respectively.

Scaling

A third special internal quantity called 'scale' is used to determine the scale of calculated quantities. Numbers may have up to 99 decimal digits after the decimal point. This fractional part is retained in further computations. We refer to the number of digits after the decimal point of a number as its scale.

When two scaled numbers are combined by means of one of the arithmetic operations, the result has a scale determined by the following rules. For addition and subtraction, the scale of the result is the larger of the scales of the two operands. In this case, there is never any truncation of the result. For multiplications, the scale of the result is never less than the maximum of the two scales of the operands, never more than the sum of the scales of the operands and, subject to those two restrictions, the scale of the result is set equal to the contents of the internal quantity 'scale'. The scale of a quotient is the contents of the internal quantity 'scale'. The scale of a remainder is the sum of the scales of the quotient and the divisor. The result of an exponentiation is scaled as if the implied multiplications were performed. An exponent must be an integer. The scale of a square root is set to the maximum of the scale of the argument and the contents of 'scale'.

All of the internal operations are actually carried out in terms of integers, with digits being discarded when necessary. In every case where digits are discarded, truncation and not rounding is performed.

The contents of 'scale' must be no greater than 99 and no less than 0. It is initially set to 0. In case you need more than 99 fraction digits, you may arrange your own scaling.

The internal quantities 'scale', 'ibase', and 'obase' can be used in expressions just like other variables. The line

```
scale = scale + 1
```

increases the value of 'scale' by one, and the line

```
scale
```

causes the current value of 'scale' to be printed.

The value of 'scale' retains its meaning as a number of decimal digits to be retained in internal computation even when 'ibase' or 'obase' are not equal to 10. The internal computations (which are still conducted in decimal, regardless of the bases) are performed to the specified number of decimal digits, never hexadecimal or octal or any other kind of digits.

Functions

The name of a function is a single lower-case letter. Function names are permitted to collide with simple variable names. Twenty-six different defined functions are permitted in addition to the twenty-six variable names. The line

```
define a(x){
```

begins the definition of a function with one argument. This line must be followed by one or more statements, which make up the body of the function, ending with a right brace }. Return of control from a function occurs when a return statement is executed or when the end of the function is reached. The return statement can take either of the two forms

```
return  
return(x)
```

In the first case, the value of the function is 0, and in the second, the value of the expression in parentheses.

Variables used in the function can be declared as automatic by a statement of the form

```
auto x,y,z
```

There can be only one 'auto' statement in a function and it must be the first statement in the definition. These automatic variables are allocated space and initialized to zero on entry to the function and thrown away on return. The values of any variables with the same names outside the function are not disturbed. Functions may be called recursively and the automatic variables at each level of call are protected. The parameters named in a function definition are treated in the same way as the automatic variables of that function with the single exception that they are given a value on entry to the function. An example of a function definition is

```
define a(x,y){  
    auto z  
    z = x*y  
    return(z)  
}
```

The value of this function, when called, will be the product of its two arguments.

A function is called by the appearance of its name followed by a string of arguments enclosed in parentheses and separated by commas. The result is unpredictable if the wrong number of arguments is used.

Functions with no arguments are defined and called using parentheses with nothing between them: b().

If the function *a* above has been defined, then the line

```
a(7,3.14)
```

would cause the result 21.98 to be printed and the line

```
x = a(a(3,4),5)
```

would cause the value of *x* to become 60.

Subscripted Variables

A single lower-case letter variable name followed by an expression in brackets is called a subscripted variable (an array element). The variable name is called the array name and the expression in brackets is called the subscript. Only one-dimensional arrays are permitted. The names of arrays are permitted to collide with the names of simple variables and function names. Any fractional part of a subscript is discarded before use. Subscripts must be greater than or equal to zero and less than or equal to 2047.

Subscripted variables may be freely used in expressions, in function calls, and in return statements.

An array name may be used as an argument to a function, or may be declared as automatic in a function definition by the use of empty brackets:

```
f(a[])  
define f(a[])  
auto a[]
```

When an array name is so used, the whole contents of the array are copied for the use of the function, and thrown away on exit from the function. Array names which refer to whole arrays cannot be used in any other contexts.

Control Statements

The 'if', the 'while', and the 'for' statements may be used to alter the flow within programs or to cause iteration. The range of each of them is a statement or a compound statement consisting of a collection of statements enclosed in braces. They are written in the following way

```
if(relation) statement  
while(relation) statement  
for(expression1; relation; expression2) statement
```

or

```
if(relation) {statements}  
while(relation) {statements}  
for(expression1; relation; expression2) {statements}
```

A relation in one of the control statements is an expression of the form

$x > y$

where two expressions are related by one of the six relational operators $<$, $>$, $<=$, $>=$, $==$, or $!=$. The relation $==$ stands for 'equal to' and $!=$ stands for 'not equal to'. The meaning of the remaining relational operators is clear.

BEWARE of using $=$ instead of $==$ in a relational. Unfortunately, both of them are legal, so you will not get a diagnostic message, but $=$ really will not do a comparison.

The 'if' statement causes execution of its range if and only if the relation is true. Then control passes to the next statement in sequence.

The 'while' statement causes execution of its range repeatedly as long as the relation is true. The relation is tested before each execution of its range and if the relation is false, control passes to the next statement beyond the range of the while.

The 'for' statement begins by executing 'expression1'. Then the relation is tested and, if true, the statements in the range of the 'for' are executed. Then 'expression2' is executed. The relation is tested, and so on. The typical use of the 'for' statement is for a controlled iteration, as in the statement

```
for(i=1; i<=10; i=i+1) i
```

which will print the integers from 1 to 10. Here are some examples of the use of the control statements.

```
define f(n){
  auto i, x
  x=1
  for(i=1; i<=n; i=i+1) x=x*i
  return(x)
}
```

The line

```
f(a)
```

will print a factorial if a is a positive integer. Here is the definition of a function which will compute values of the binomial coefficient (m and n are assumed to be positive integers).

```
define b(n,m){
  auto x, j
  x=1
  for(j=1; j<=m; j=j+1) x=x*(n-j+1)/j
  return(x)
}
```

The following function computes values of the exponential function by summing the appropriate series without regard for possible truncation errors:

```
scale = 20
define e(x){
  auto a, b, c, d, n
  a = 1
  b = 1
  c = 1
  d = 0
  n = 1
  while(1==1){
    a = a*x
    b = b*n
    c = c + a/b
    n = n + 1
    if(c==d) return(c)
    d = c
  }
}
```

Some Details

There are some language features that every user should know about even if he will not use them.

Normally statements are typed one to a line. It is also permissible to type several statements on a line separated by semicolons.

If an assignment statement is parenthesized, it then has a value and it can be used anywhere that an expression can. For example, the line

(x=y+17)

not only makes the indicated assignment, but also prints the resulting value.

Here is an example of a use of the value of an assignment statement even when it is not parenthesized.

x = a[i=i+1]

causes a value to be assigned to x and also increments i before it is used as a subscript.

The following constructs work in BC in exactly the same manner as they do in the C language. Consult the appendix or the C manuals [2,3] for their exact workings.

x=y=z	is the same as	x=(y=z)
x =+ y		x = x+y
x =- y		x = x-y
x =* y		x = x*y
x =/ y		x = x/y
x =% y		x = x%y
x =^ y		x = x^y
x++		(x=x+1)-1
x--		(x=x-1)+1
++x		x = x+1
--x		x = x-1

Even if you don't intend to use the constructs, if you type one inadvertently, something correct but unexpected may happen.

WARNING! In some of these constructions, spaces are significant. There is a real difference between x=-y and x= -y. The first replaces x by x-y and the second by -y.

Three Important Things

1. To exit a BC program, type 'quit'.
2. There is a comment convention identical to that of C and of PL/I. Comments begin with '/' and end with '*'.
3. There is a library of math functions which may be obtained by typing at command level

bc -l

This command will load a set of library functions which, at the time of writing, consists of sine (named 's'), cosine ('c'), arctangent ('a'), natural logarithm ('l'), exponential ('e') and Bessel functions of integer order ('j(n,x)'). Doubtless more functions will be added in time. The library sets the scale to 20. You can reset it to something else if you like. The design of these mathematical library routines is discussed elsewhere [4].

If you type

bc file ...

BC will read and execute the named file or files before accepting commands from the keyboard. In this way, you may load your favorite programs and function definitions.

Acknowledgement

The compiler is written in YACC [5]; its original version was written by S. C. Johnson.

References

- [1] K. Thompson and D. M. Ritchie, *UNIX Programmer's Manual*, Fifth Edition (1974)
- [2] D. M. Ritchie, *C Reference Manual*.
- [3] B. W. Kernighan, *Programming in C: A Tutorial*.
- [4] Robert Morris, *A Library of Reference Standard Mathematical Subroutines*, Internal memorandum, Bell Laboratories, 1975.
- [5] S. C. Johnson, *YACC, Yet Another Compiler-Compiler*.
- [6] R. Morris and L. L. Cherry, *DC - An Interactive Desk Calculator*.

Appendix

1. Notation

In the following pages syntactic categories are in *italics*; literals are in **bold**; material in brackets [] is optional.

2. Tokens

Tokens consist of keywords, identifiers, constants, operators, and separators. Token separators may be blanks, tabs or comments. Newline characters or semicolons separate statements.

2.1. Comments

Comments are introduced by the characters /* and terminated by */.

2.2. Identifiers

There are three kinds of identifiers – ordinary identifiers, array identifiers and function identifiers. All three types consist of single lower-case letters. Array identifiers are followed by square brackets, possibly enclosing an expression describing a subscript. Arrays are singly dimensioned and may contain up to 2048 elements. Indexing begins at zero so an array may be indexed from 0 to 2047. Subscripts are truncated to integers. Function identifiers are followed by parentheses, possibly enclosing arguments. The three types of identifiers do not conflict; a program can have a variable named *x*, an array named *x* and a function named *x*, all of which are separate and distinct.

2.3. Keywords

The following are reserved keywords:

ibase	if
obase	break
scale	define
sqrt	auto
length	return
while	quit
for	

2.4. Constants

Constants consist of arbitrarily long numbers with an optional decimal point. The hexadecimal digits A–F are also recognized as digits with values 10–15, respectively.

3. Expressions

The value of an expression is printed unless the main operator is an assignment. Precedence is the same as the order of presentation here, with highest appearing first. Left or right associativity, where applicable, is discussed with each operator.

3.1. Primitive expressions

3.1.1. Named expressions

Named expressions are places where values are stored. Simply stated, named expressions are legal on the left side of an assignment. The value of a named expression is the value stored in the place named.

3.1.1.1. *identifiers*

Simple identifiers are named expressions. They have an initial value of zero.

3.1.1.2. *array-name[expression]*

Array elements are named expressions. They have an initial value of zero.

3.1.1.3. *scale, ibase and obase*

The internal registers *scale*, *ibase* and *obase* are all named expressions. *scale* is the number of digits after the decimal point to be retained in arithmetic operations. *scale* has an initial value of zero. *ibase* and *obase* are the input and output number radix respectively. Both *ibase* and *obase* have initial values of 10.

3.1.2. Function calls

3.1.2.1. *function-name([expression[, expression...]])*

A function call consists of a function name followed by parentheses containing a comma-separated list of expressions, which are the function arguments. A whole array passed as an argument is specified by the array name followed by empty square brackets. All function arguments are passed by value. As a result, changes made to the formal parameters have no effect on the actual arguments. If the function terminates by executing a return statement, the value of the function is the value of the expression in the parentheses of the return statement or is zero if no expression is provided or if there is no return statement.

3.1.2.2. *sqrt(expression)*

The result is the square root of the expression. The result is truncated in the least significant decimal place. The scale of the result is the scale of the expression or the value of *scale*, whichever is larger.

3.1.2.3. *length(expression)*

The result is the total number of significant decimal digits in the expression. The scale of the result is zero.

3.1.2.4. *scale(expression)*

The result is the scale of the expression. The scale of the result is zero.

3.1.3. Constants

Constants are primitive expressions.

3.1.4. Parentheses

An expression surrounded by parentheses is a primitive expression. The parentheses are used to alter the normal precedence.

3.2. Unary operators

The unary operators bind right to left.

3.2.1. $-$ *expression*

The result is the negative of the expression.

3.2.2. $++$ *named-expression*

The named expression is incremented by one. The result is the value of the named expression after incrementing.

3.2.3. $--$ *named-expression*

The named expression is decremented by one. The result is the value of the named expression after decrementing.

3.2.4. *named-expression* $++$

The named expression is incremented by one. The result is the value of the named expression before incrementing.

3.2.5. *named-expression* $--$

The named expression is decremented by one. The result is the value of the named expression before decrementing.

3.3. Exponentiation operator

The exponentiation operator binds right to left.

3.3.1. *expression* $^$ *expression*

The result is the first expression raised to the power of the second expression. The second expression must be an integer. If a is the scale of the left expression and b is the absolute value of the right expression, then the scale of the result is:

$$\min(a \times b, \max(\text{scale}, a))$$

3.4. Multiplicative operators

The operators $*$, $/$, $\%$ bind left to right.

3.4.1. *expression* $*$ *expression*

The result is the product of the two expressions. If a and b are the scales of the two expressions, then the scale of the result is:

$$\min(a + b, \max(\text{scale}, a, b))$$

3.4.2. *expression / expression*

The result is the quotient of the two expressions. The scale of the result is the value of scale.

3.4.3. *expression % expression*

The % operator produces the remainder of the division of the two expressions. More precisely, $a\%b$ is $a - a/b * b$.

The scale of the result is the sum of the scale of the divisor and the value of scale

3.5. Additive operators

The additive operators bind left to right.

3.5.1. *expression + expression*

The result is the sum of the two expressions. The scale of the result is the maximum of the scales of the expressions.

3.5.2. *expression - expression*

The result is the difference of the two expressions. The scale of the result is the maximum of the scales of the expressions.

3.6. assignment operators

The assignment operators bind right to left.

3.6.1. *named-expression = expression*

This expression results in assigning the value of the expression on the right to the named expression on the left.

3.6.2. *named-expression =+ expression*

3.6.3. *named-expression =- expression*

3.6.4. *named-expression =* expression*

3.6.5. *named-expression =/ expression*

3.6.6. *named-expression =% expression*

3.6.7. *named-expression ^= expression*

The result of the above expressions is equivalent to "named expression = named expression OP expression", where OP is the operator after the = sign.

4. Relations

Unlike all other operators, the relational operators are only valid as the object of an if, while, or inside a for statement.

4.1. *expression < expression*

4.2. *expression > expression*

4.3. *expression <= expression*

4.4. *expression >= expression*

4.5. *expression == expression*

4.6. *expression != expression*

5. Storage classes

There are only two storage classes in BC, global and automatic (local). Only identifiers that are to be local to a function need be declared with the **auto** command. The arguments to a function are local to the function. All other identifiers are assumed to be global and available to all functions. All identifiers, global and local, have initial values of zero. Identifiers declared as **auto** are allocated on entry to the function and released on returning from the function. They therefore do not retain values between function calls. **auto** arrays are specified by the array name followed by empty square brackets.

Automatic variables in BC do not work in exactly the same way as in either C or PL/I. On entry to a function, the old values of the names that appear as parameters and as automatic variables are pushed onto a stack. Until return is made from the function, reference to these names refers only to the new values.

6. Statements

Statements must be separated by semicolon or newline. Except where altered by control statements, execution is sequential.

6.1. Expression statements

When a statement is an expression, unless the main operator is an assignment, the value of the expression is printed, followed by a newline character.

6.2. Compound statements

Statements may be grouped together and used when one statement is expected by surrounding them with { }.

6.3. Quoted string statements

"any string"

This statement prints the string inside the quotes.

6.4. If statements

if (relation) statement

The substatement is executed if the relation is true.

6.5. While statements

while(*relation*) *statement*

The statement is executed while the relation is true. The test occurs before each execution of the statement.

6.6. For statements

for(*expression*; *relation*; *expression*) *statement*

The for statement is the same as

```
first-expression  
while(relation) {  
    statement  
    last-expression  
}
```

All three expressions must be present.

6.7. Break statements

break

break causes termination of a **for** or **while** statement.

6.8. Auto statements

auto *identifier* [, *identifier*]

The auto statement causes the values of the identifiers to be pushed down. The identifiers can be ordinary identifiers or array identifiers. Array identifiers are specified by following the array name by empty square brackets. The auto statement must be the first statement in a function definition.

6.9. Define statements

```
define( [parameter [ , parameter ... ] ] ) {  
    statements  
}
```

The define statement defines a function. The parameters may be ordinary identifiers or array names. Array names must be followed by empty square brackets.

6.10. Return statements

return

return(*expression*)

The return statement causes termination of a function, popping of its auto variables, and specifies the result of the function. The first form is equivalent to **return**(0). The result of the function is the result of the expression in parentheses.

6.11. Quit

The quit statement stops execution of a BC program and returns control to UNIX when it is first encountered. Because it is not treated as an executable statement, it cannot be used in a function definition or in an **if**, **for**, or **while** statement.

DC – An Interactive Desk Calculator

Robert Morris

Lorinda Cherry

Bell Laboratories,
Murray Hill, New Jersey 07974

DC is an arbitrary precision arithmetic package implemented on the UNIX time-sharing system in the form of an interactive desk calculator. It works like a stacking calculator using reverse Polish notation. Ordinarily DC operates on decimal integers, but one may specify an input base, output base, and a number of fractional digits to be maintained.

A language called BC [1] has been developed which accepts programs written in the familiar style of higher-level programming languages and compiles output which is interpreted by DC. Some of the commands described below were designed for the compiler interface and are not easy for a human user to manipulate.

Numbers that are typed into DC are put on a push-down stack. DC commands work by taking the top number or two off the stack, performing the desired operation, and pushing the result on the stack. If an argument is given, input is taken from that file until its end, then from the standard input.

SYNOPTIC DESCRIPTION

Here we describe the DC commands that are intended for use by people. The additional commands that are intended to be invoked by compiled output are described in the detailed description.

Any number of commands are permitted on a line. Blanks and new-line characters are ignored except within numbers and in places where a register name is expected.

The following constructions are recognized:

number

The value of the number is pushed onto the main stack. A number is an unbroken string of the digits 0-9 and the capital letters A-F which are treated as digits with values 10-15 respectively. The number may be preceded by an underscore to input a negative number. Numbers may contain decimal points.

+ - * % ^

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. The result of a division is an integer truncated toward zero. See the detailed description below for the treatment of numbers with decimal points. An exponent must not have any digits after the decimal point.

sx

The top of the main stack is popped and stored into a register named *x*, where *x* may be any character. If the *s* is capitalized, *x* is treated as a stack and the value is pushed onto it. Any character, even blank or new-line, is a valid register name.

lx

The value in register *x* is pushed onto the stack. The register *x* is not altered. If the *l* is capitalized, register *x* is treated as a stack and its top value is popped onto the main stack.

All registers start with empty value which is treated as a zero by the command *l* and is treated as an error by the command *L*.

d

The top value on the stack 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.

x

treats the top element of the stack as a character string, removes it from the stack, and executes it as a string of DC commands.

[...]

puts the bracketed character string onto the top of the stack.

q

exits the program. If executing a string, the recursion level is popped by two. If *q* is capitalized, the top value on the stack is popped and the string execution level is popped by that value.

<x >x =x !<x !>x !=x

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

v

replaces the top element on the stack by its square root. The square root of an integer is truncated to an integer. For the treatment of numbers with decimal points, see the detailed description below.

- !
- interprets the rest of the line as a UNIX command. Control returns to DC when the UNIX command terminates.
- c
- All values on the stack are popped; the stack becomes empty.
- i
- The top value on the stack is popped and used as the number radix for further input. If i is capitalized, the value of the input base is pushed onto the stack. No mechanism has been provided for the input of arbitrary numbers in bases less than 1 or greater than 16.
- o
- The top value on the stack is popped and used as the number radix for further output. If o is capitalized, the value of the output base is pushed onto the stack.
- k
- The top of the stack is popped, and that value is used as a scale factor that influences the number of decimal places that are maintained during multiplication, division, and exponentiation. The scale factor must be greater than or equal to zero and less than 100. If k is capitalized, the value of the scale factor is pushed onto the stack.
- z
- The value of the stack level is pushed onto the stack.
- ?
- A line of input is taken from the input source (usually the console) and executed.

DETAILED DESCRIPTION

Internal Representation of Numbers

Numbers are stored internally using a dynamic storage allocator. Numbers are kept in the form of a string of digits to the base 100 stored one digit per byte (centennial digits). The string is stored with the low-order digit at the beginning of the string. For example, the representation of 157 is 57,1. After any arithmetic operation on a number, care is taken that all digits are in the range 0-99 and that the number has no leading zeros. The number zero is represented by the empty string.

Negative numbers are represented in the 100's complement notation, which is analogous to two's complement notation for binary numbers. The high order digit of a negative number is always -1 and all other digits are in the range 0-99. The digit preceding the high order -1 digit is never a 99. The representation of -157 is 43,98,-1. We shall call this the canonical form of a number. The advantage of this kind of representation of negative numbers is ease of addition. When addition is performed digit by digit, the result is formally correct. The result need only be modified, if necessary, to put it into canonical form.

Because the largest valid digit is 99 and the byte can hold numbers twice that large, addition can be carried out and the handling of carries done later when that is convenient, as it sometimes is.

An additional byte is stored with each number beyond the high order digit to indicate the number of assumed decimal digits after the decimal point. The representation of .001 is 1,3

where the scale has been italicized to emphasize the fact that it is not the high order digit. The value of this extra byte is called the **scale factor** of the number.

The Allocator

DC uses a dynamic string storage allocator for all of its internal storage. All reading and writing of numbers internally is done through the allocator. Associated with each string in the allocator is a four-word header containing pointers to the beginning of the string, the end of the string, the next place to write, and the next place to read. Communication between the allocator and DC is done via pointers to these headers.

The allocator initially has one large string on a list of free strings. All headers except the one pointing to this string are on a list of free headers. Requests for strings are made by size. The size of the string actually supplied is the next higher power of 2. When a request for a string is made, the allocator first checks the free list to see if there is a string of the desired size. If none is found, the allocator finds the next larger free string and splits it repeatedly until it has a string of the right size. Left-over strings are put on the free list. If there are no larger strings, the allocator tries to coalesce smaller free strings into larger ones. Since all strings are the result of splitting large strings, each string has a neighbor that is next to it in core and, if free, can be combined with it to make a string twice as long. This is an implementation of the 'buddy system' of allocation described in [2].

Failing to find a string of the proper length after coalescing, the allocator asks the system for more space. The amount of space on the system is the only limitation on the size and number of strings in DC. If at any time in the process of trying to allocate a string, the allocator runs out of headers, it also asks the system for more space.

There are routines in the allocator for reading, writing, copying, rewinding, forward-spacing, and backspacing strings. All string manipulation is done using these routines.

The reading and writing routines increment the read pointer or write pointer so that the characters of a string are read or written in succession by a series of read or write calls. The write pointer is interpreted as the end of the information-containing portion of a string and a call to read beyond that point returns an end-of-string indication. An attempt to write beyond the end of a string causes the allocator to allocate a larger space and then copy the old string into the larger block.

Internal Arithmetic

All arithmetic operations are done on integers. The operands (or operand) needed for the operation are popped from the main stack and their scale factors stripped off. Zeros are added or digits removed as necessary to get a properly scaled result from the internal arithmetic routine. For example, if the scale of the operands is different and decimal alignment is required, as it is for addition, zeros are appended to the operand with the smaller scale. After performing the required arithmetic operation, the proper scale factor is appended to the end of the number before it is pushed on the stack.

A register called **scale** plays a part in the results of most arithmetic operations. **scale** is the bound on the number of decimal places retained in arithmetic computations. **scale** may be set to the number on the top of the stack truncated to an integer with the **k** command. **K** may be used to push the value of **scale** on the stack. **scale** must be greater than or equal to 0 and less than 100. The descriptions of the individual arithmetic operations will include the exact effect of **scale** on the computations.

Addition and Subtraction

The scales of the two numbers are compared and trailing zeros are supplied to the number with the lower scale to give both numbers the same scale. The number with the smaller scale is multiplied by 10 if the difference of the scales is odd. The scale of the result is then set to the larger of the scales of the two operands.

Subtraction is performed by negating the number to be subtracted and proceeding as in addition.

Finally, the addition is performed digit by digit from the low order end of the number. The carries are propagated in the usual way. The resulting number is brought into canonical form, which may require stripping of leading zeros, or for negative numbers replacing the high-order configuration 99,-1 by the digit -1. In any case, digits which are not in the range 0-99 must be brought into that range, propagating any carries or borrows that result.

Multiplication

The scales are removed from the two operands and saved. The operands are both made positive. Then multiplication is performed in a digit by digit manner that exactly mimics the hand method of multiplying. The first number is multiplied by each digit of the second number, beginning with its low order digit. The intermediate products are accumulated into a partial sum which becomes the final product. The product is put into the canonical form and its sign is computed from the signs of the original operands.

The scale of the result is set equal to the sum of the scales of the two operands. If that scale is larger than the internal register scale and also larger than both of the scales of the two operands, then the scale of the result is set equal to the largest of these three last quantities.

Division

The scales are removed from the two operands. Zeros are appended or digits removed from the dividend to make the scale of the result of the integer division equal to the internal quantity scale. The signs are removed and saved.

Division is performed much as it would be done by hand. The difference of the lengths of the two numbers is computed. If the divisor is longer than the dividend, zero is returned. Otherwise the top digit of the divisor is divided into the top two digits of the dividend. The result is used as the first (high-order) digit of the quotient. It may turn out to be one unit too low, but if it is, the next trial quotient will be larger than 99 and this will be adjusted at the end of the process. The trial digit is multiplied by the divisor and the result subtracted from the dividend and the process is repeated to get additional quotient digits until the remaining dividend is smaller than the divisor. At the end, the digits of the quotient are put into the canonical form, with propagation of carry as needed. The sign is set from the sign of the operands.

Remainder

The division routine is called and division is performed exactly as described. The quantity returned is the remains of the dividend at the end of the divide process. Since division truncates toward zero, remainders have the same sign as the dividend. The scale of the remainder is set to the maximum of the scale of the dividend and the scale of the quotient plus the scale of the divisor.

Square Root

The scale is stripped from the operand. Zeros are added if necessary to make the integer result have a scale that is the larger of the internal quantity scale and the scale of the operand.

The method used to compute $\text{sqrt}(y)$ is Newton's method with successive approximations by the rule

$$x_{n+1} = \frac{1}{2} \left(x_n + \frac{y}{x_n} \right)$$

The initial guess is found by taking the integer square root of the top two digits.

Exponentiation

Only exponents with zero scale factor are handled. If the exponent is zero, then the result is 1. If the exponent is negative, then it is made positive and the base is divided into one. The scale of the base is removed.

The integer exponent is viewed as a binary number. The base is repeatedly squared and the result is obtained as a product of those powers of the base that correspond to the positions of the one-bits in the binary representation of the exponent. Enough digits of the result removed to make the scale of the result the same as if the indicated multiplication had been performed.

Input Conversion and Base

Numbers are converted to the internal representation as they are read in. The scale stored with a number is simply the number of fractional digits input. Negative numbers are indicated by preceding the number with a `_`. The hexadecimal digits A-F correspond to the numbers 10-15 regardless of input base. The `i` command can be used to change the base of the input numbers. This command pops the stack, truncates the resulting number to an integer, and uses it as the input base for all further input. The input base is initialized to 10 but may, for example be changed to 8 or 16 to do octal or hexadecimal to decimal conversions. The command `I` will push the value of the input base on the stack.

Output Commands

The command `p` causes the top of the stack to be printed. It does not remove the top of the stack. All of the stack and internal registers can be output by typing the command `f`. The `o` command can be used to change the output base. This command uses the top of the stack, truncated to an integer as the base for all further output. The output base is initialized to 10. It will work correctly for any base. The command `O` pushes the value of the output base on the stack.

Output Format and Base

The input and output bases only affect the interpretation of numbers on input and output; they have no effect on arithmetic computations. Large numbers are output with 70 characters per line; a `\` indicates a continued line. All choices of input and output bases work correctly, although not all are useful. A particularly useful output base is 100000, which has the effect of grouping digits in fives. Bases of 8 and 16 can be used for decimal-octal or decimal-hexadecimal conversions.

Internal Registers

Numbers or strings may be stored in internal registers or loaded on the stack from registers with the commands `s` and `l`. The command `sx` pops the top of the stack and stores the result in register `x`. `x` can be any character. `lx` puts the contents of register `x` on the top of the stack. The `l` command has no effect on the contents of register `x`. The `s` command, however, is destructive.

Stack Commands

The command `c` clears the stack. The command `d` pushes a duplicate of the number on the top of the stack on the stack. The command `z` pushes the stack size on the stack. The command `X` replaces the number on the top of the stack with its scale factor. The command `Z` replaces the top of the stack with its length.

Subroutine Definitions and Calls

Enclosing a string in `||` pushes the ascii string on the stack. The `q` command quits or in executing a string, pops the recursion levels by two.

Internal Registers — Programming DC

The load and store commands together with `||` to store strings, `x` to execute and the testing commands '`<`', '`>`', '`=`', '`!<`', '`!>`', '`!=`' can be used to program DC. The `x` command assumes the top of the stack is a string of DC commands and executes it. The testing commands compare the top two elements on the stack and if the relation holds, execute the register that follows the relation. For example, to print the numbers 0-9,

```
[lipl+ si li10>a]sa
0si lax
```

Push-Down Registers and Arrays

These commands were designed for use by a compiler, not by people. They involve push-down registers and arrays. In addition to the stack that commands work on, DC can be thought of as having individual stacks for each register. These registers are operated on by the commands `S` and `L`. `Sx` pushes the top value of the main stack onto the stack for the register `x`. `Lx` pops the stack for register `x` and puts the result on the main stack. The commands `s` and `l` also work on registers but not as push-down stacks. `l` doesn't effect the top of the register stack, and `s` destroys what was there before.

The commands to work on arrays are `:` and `;`. `:x` pops the stack and uses this value as an index into the array `x`. The next element on the stack is stored at this index in `x`. An index must be greater than or equal to 0 and less than 2048. `;x` is the command to load the main stack from the array `x`. The value on the top of the stack is the index into the array `x` of the value to be loaded.

Miscellaneous Commands

The command `!` interprets the rest of the line as a UNIX command and passes it to UNIX to execute. One other compiler command is `Q`. This command uses the top of the stack as the number of levels of recursion to skip.

DESIGN CHOICES

The real reason for the use of a dynamic storage allocator was that a general purpose program could be (and in fact has been) used for a variety of other tasks. The allocator has some value for input and for compiling (i.e. the bracket [...] commands) where it cannot be known in advance how long a string will be. The result was that at a modest cost in execution time, all considerations of string allocation and sizes of strings were removed from the remainder of the program and debugging was made easier. The allocation method used wastes approximately 25% of available space.

The choice of 100 as a base for internal arithmetic seemingly has no compelling advantage. Yet the base cannot exceed 127 because of hardware limitations and at the cost of 5% in space, debugging was made a great deal easier and decimal output was made much faster.

The reason for a stack-type arithmetic design was to permit all DC commands from addition to subroutine execution to be implemented in essentially the same way. The result was a considerable degree of logical separation of the final program into modules with very little communication between modules.

The rationale for the lack of interaction between the scale and the bases was to provide an understandable means of proceeding after a change of base or scale when numbers had already been entered. An earlier implementation which had global notions of scale and base did not work out well. If the value of *scale* were to be interpreted in the current input or output base, then a change of base or scale in the midst of a computation would cause great confusion in the interpretation of the results. The current scheme has the advantage that the value of the input and output bases are only used for input and output, respectively, and they are ignored in all other operations. The value of *scale* is not used for any essential purpose by any part of the program and it is used only to prevent the number of decimal places resulting from the arithmetic operations from growing beyond all bounds.


The design rationale for the choices for the scales of the results of arithmetic were that in no case should any significant digits be thrown away if, on appearances, the user actually wanted them. Thus, if the user wants to add the numbers 1.5 and 3.517, it seemed reasonable to give him the result 5.017 without requiring him to unnecessarily specify his rather obvious requirements for precision.


On the the other hand, multiplication and exponentiation produce results with many more digits than their operands and it seemed reasonable to give as a minimum the number of decimal places in the operands but not to give more than that number of digits unless the user asked for them by specifying a value for *scale*. Square root can be handled in just the same way as multiplication. The operation of division gives arbitrarily many decimal places and there is simply no way to guess how many places the user wants. In this case only, the user must specify a *scale* to get any decimal places at all.

The scale of remainder was chosen to make it possible to recreate the dividend from the quotient and remainder. This is easy to implement; no digits are thrown away.

References

- [1] L. L. Cherry, R. Morris, *BC - An Arbitrary Precision Desk-Calculator Language*.
- [2] K. C. Knowlton, *A Fast Storage Allocator*, Comm. ACM 8, pp. 623-625 (Oct. 1965)





YACC — Yet Another Compiler-Compiler

Stephen C. Johnson

Bell Laboratories,
Murray Hill, New Jersey 07974

ABSTRACT

Computer program input generally has some structure; in fact, every computer program which does input can be thought of as defining an "input language" which it accepts. The input languages may be as complex as a programming language, or as simple as a sequence of numbers. Unfortunately, standard input facilities are restricted, difficult to use and change, and do not completely check their inputs for validity.

Yacc provides a general tool for controlling the input to a computer program. The Yacc user describes the structures of his input, together with code which is to be invoked when each such structure is recognized. Yacc turns such a specification into a subroutine which may be invoked to handle the input process; frequently, it is convenient and appropriate to have most of the flow of control in the user's application handled by this subroutine.

The input subroutine produced by Yacc calls a user supplied routine to return the next basic input item. Thus, the user can specify his input in terms of individual input characters, or, if he wishes, in terms of higher level constructs such as names and numbers. The user supplied routine may also handle idiomatic features such as comment and continuation conventions, which typically defy easy specification.

Yacc is written in C[7], and runs under UNIX. The subroutine which is output may be in C or in Ratfor[4], at the user's choice; Ratfor permits translation of the output subroutine into portable Fortran[5]. The class of specifications accepted is a very general one, called LALR(1) grammars with disambiguating rules. The theory behind Yacc has been described elsewhere[1,2,3].

Yacc was originally designed to help produce the "front end" of compilers; in addition to this use, it has been successfully used in many application programs, including a phototypesetter language, a document retrieval system, a Fortran debugging system, and the Ratfor compiler.



YACC — Yet Another Compiler-Compiler

Stephen C. Johnson

Bell Laboratories,
Murray Hill, New Jersey 07974

Section 0: Introduction

Yacc provides a general tool for imposing structure on the input to a computer program. The Yacc user prepares a specification of the input process; this includes rules which describe the input structure, code which is to be invoked when these structures are recognized, and a low-level routine to do the basic input. Yacc then produces a subroutine to do the input procedure; this subroutine, called a *parser*, calls the user-supplied low-level input routine (called the *lexical analyzer*) to pick up the basic items (called *tokens*) from the input stream. These tokens are organized according to the input structure rules, called *grammar rules*; when one of these rules has been recognized, then the user code supplied for this rule, called an *action*, is invoked; actions have the ability to return values and make use of the values of other actions.

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

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

Here, `date`, `month_name`, `day`, and `year` represent structures of interest in the input process; presumably, `month_name`, `day`, and `year` are defined elsewhere. The comma “,” is quoted by single quotes; this implies that the comma is to appear literally in the input. The colon and semicolon merely serve as punctuation in the rule, and have no significance in controlling the input. Thus, with proper definitions, the input

```
July 4, 1776
```

might be matched by the above rule.

As we mentioned above, an important part of the input process is carried out by the lexical analyzer. This user routine reads the true input stream, recognizing those structures which are more conveniently or more efficiently recognized directly, and communicates these recognized tokens to the parser. For historical reasons, the name of a structure recognized by the lexical analyzer is called a *terminal symbol* name, while the name of a structure recognized by the parser is called a *nonterminal symbol* name. To avoid the obvious confusion of terminology, we shall usually refer to terminal symbol names as *token names*.

There is considerable leeway in deciding whether to recognize structures by the lexical analyzer or by a grammar rule. Thus, in the above example it would be possible to have other rules of the form

```
month_name : 'J' 'a' 'n' ;  
month_name : 'F' 'e' 'b' ;
```

```
...  
month_name : 'D' 'e' 'c' ;
```

Here, the lexical analyzer would only need to recognize individual letters, and `month_name` would be a nonterminal symbol. Rules of this sort tend to be a bit wasteful of time and space, and may even restrict the power of the input process (although they are easy to write). For a

more efficient input process, the lexical analyzer itself might recognize the month names, and return an indication that a month_name was seen; in this case, month_name would be a token.

Literal characters, such as “,”, must also be passed through the lexical analyzer, and are considered tokens.

As an example of the flexibility of the grammar rule approach, we might add to the above specifications the rule

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

and thus optionally allow the form

```
7/4/1776
```

as a synonym for

```
July 4, 1776
```

In most cases, this new rule could be “slipped in” to a working system with minimal effort, and a very small chance of disrupting existing input.

Frequently, the input being read does not conform to the specifications due to errors in the input. The parsers produced by Yacc have the very desirable property that they will detect these input errors at the earliest place at which this can be done with a left-to-right scan; thus, not only is the chance of reading and computing with bad input data substantially reduced, but the bad data can usually be quickly found. Error handling facilities, entered as part of the input specifications, frequently permit the reentry of bad data, or the continuation of the input process after skipping over the bad data.

In some cases, Yacc fails to produce a parser when given a set of specifications. For example, the specifications may be self contradictory, or they may require a more powerful recognition mechanism than that available to Yacc. The former cases probably represent true design errors; the latter cases can often be corrected by making the lexical analyzer more powerful, or by rewriting some of the grammar rules. The class of specifications which Yacc can handle compares very favorably with other systems of this type; moreover, the constructions which are difficult for Yacc to handle are also frequently difficult for human beings to handle. Some users have reported that the discipline of formulating valid Yacc specifications for their input revealed errors of conception or design early in the program development.

The next several sections describe the basic process of preparing a Yacc specification; Section 1 describes the preparation of grammar rules, Section 2 the preparation of the user supplied actions associated with these rules, and Section 3 the preparation of lexical analyzers. In Section 4, we discuss the diagnostics produced when Yacc is unable to produce a parser from the given specifications. This section also describes a simple, frequently useful mechanism for handling operator precedences. Section 5 discusses error detection and recovery. Sections 6C and 6R discuss the operating environment and special features of the subroutines which Yacc produces in C and Ratfor, respectively. Section 7 gives some hints which may lead to better designed, more efficient, and clearer specifications. Finally, Section 8 has a brief summary. Appendix A has a brief example, and Appendix B tells how to run Yacc on the UNIX operating system. Appendix C has a brief description of mechanisms and syntax which are no longer actively supported, but which are provided for historical continuity with older versions of Yacc.

Section 1: Basic Specifications

As we noted above, names refer to either tokens or nonterminal symbols. Yacc requires those names which will be used as token names to be declared as such. In addition, for reasons which will be discussed in Section 3, it is usually desirable to include the lexical analyzer as part of the specification file; it may be useful to include other programs as well. Thus, every specification file consists of three sections: the *declarations*, *(grammar) rules*, and *programs*. The sections are separated by double percent “%%” marks. (The per-cent “%” is generally used in Yacc specifications as an escape character.)

In other words, a full specification file looks like

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

The declaration section may be empty. Moreover, if the programs section is omitted, the second %% mark may be omitted also; thus, the smallest legal Yacc specification is

```
%%
rules
```

Blanks, tabs, and newlines are ignored except that they may not appear in names or multi-character reserved symbols. Comments may appear wherever a name or operator is legal; they are enclosed in /* ... */, as in C and PL/I.

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

```
A : BODY ;
```

A represents a nonterminal name, and BODY represents a sequence of zero or more names and literals. Notice that the colon and the semicolon are Yacc punctuation.

Names may be of arbitrary length, and may be made up of letters, dot “.”, underscore “_”, and non-initial digits. Notice that Yacc considers that upper and lower case letters are distinct. The names used in the body of a grammar rule may represent tokens or nonterminal symbols.

A literal consists of a character enclosed in single quotes “’”. As in C, the backslash “\” is an escape character within literals, and all the C escapes are recognized. Thus

```
‘\n’ represents newline
‘\r’ represents return
‘\’  represents single quote “’”
‘\\’ represents backslash “\”
‘\t’ represents tab
‘\b’ represents backspace
‘\xxx’ represents “xxx” in octal
```

For a number of technical reasons, the nul character (‘\0’ or 000) should never be used in grammar rules.

If there are several grammar rules with the same left hand side, the vertical bar “|” can be used to avoid rewriting the left hand side. In addition, the semicolon at the end of a rule can be dropped before a vertical bar. Thus the grammar rules

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

can be given to Yacc as

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

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

If a nonterminal symbol matches the empty string, this can be indicated in the obvious way:

```
empty : ;
```

As we mentioned above, names which represent tokens must be declared as such. The simplest way of doing this is to write

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

in the declarations section. (See Sections 3 and 4 for much more discussion). Every name not defined in the declarations section is assumed to represent a nonterminal symbol. If, by the end of the rules section, some nonterminal symbol has not appeared on the left of any rule, then an error message is produced and Yacc halts.

The left hand side of the *first* grammar rule in the grammar rules section has special importance; it is taken to be the controlling nonterminal symbol for the entire input process; in technical language it is called the *start symbol*. In effect, the parser is designed to recognize the start symbol; thus, this symbol generally represents the largest, most general structure described by the grammar rules.

The end of the input is signaled by a special token, called the *endmarker*. If the tokens up to, but not including, the endmarker form a structure which matches the start symbol, the parser subroutine returns to its caller when the endmarker is seen; we say that it *accepts* the input. If the endmarker is seen in any other context, it is an error.

It is the job of the user supplied lexical analyzer to return the endmarker when appropriate; see section 3, below. Frequently, the endmarker token represents some reasonably obvious I/O status, such as "end-of-file" or "end-of-record".

Section 2: Actions

To each grammar rule, the user may associate an action to be performed each time the rule is recognized in the input process. This action may return a value, and may obtain the values returned by previous actions in the grammar rule. In addition, the lexical analyzer can return values for tokens, if desired.

When invoking Yacc, the user specifies a programming language; currently, Ratfor and C are supported. An action is an arbitrary statement in this language, and as such can do input and output, call subprograms, and alter external vectors and variables (recall that a "statement" in both C and Ratfor can be compound and do many distinct tasks). An action is specified by an equal sign "=" at the end of a grammar rule, followed by one or more statements, enclosed in curly braces "{" and "}". For example,

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

and

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

are grammar rules with actions in C. A grammar rule with an action need not end with a semicolon; in fact, it is an error to have a semicolon before the equal sign.

To facilitate easy communication between the actions and the parser, the action statements are altered slightly. The symbol "dollar sign" "\$" is used as a signal to Yacc in this context.

To return a value, the action normally sets the pseudo-variable "\$\$" to some integer value. For example, an action which does nothing but return the value 1 is

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

To obtain the values returned by previous actions and the lexical analyzer, the action may use the (integer) pseudo-variables \$1, \$2, . . . , which refer to the values returned by the components of the right side of a rule, reading from left to right. Thus, if the rule is

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

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

As a more concrete example, we might have the rule

```
expression: '(' expression '');
```

We wish the value returned by this rule to be the value of the expression in parentheses. Then we write

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

As a default, the value of a rule is the value of the first element in it (\$1). This is true even if there is no explicit action given for the rule. Thus, grammar rules of the form

```
A: B ;
```

frequently need not have an explicit action.

Notice that, although the values of actions are integers, these integers may in fact contain pointers (in C) or indices into an array (in Ratfor); in this way, actions can return and reference more complex data structures.

Sometimes, we wish to get control before a rule is fully parsed, as well as at the end of the rule. There is no explicit mechanism in Yacc to allow this; the same effect can be obtained, however, by introducing a new symbol which matches the empty string, and inserting an action for this symbol. For example, we might have a rule describing an "if" statement:

```
statement: IF '(' expr ')' THEN statement
```

Suppose that we wish to get control after seeing the right parenthesis in order to output some code. We might accomplish this by the rules:

```
statement: IF '(' expr ')' actn THEN statement  
          = { call action1 }
```

```
actn: /* matches the empty string */  
      = { call action2 }
```

Thus, the new nonterminal symbol actn matches no input, but serves only to call action2 after the right parenthesis is seen.

Frequently, it is more natural in such cases to break the rule into parts where the action is needed. Thus, the above example might also have been written

```
statement: ifpart THEN statement
          = { call action1 }

ifpart:   IF '(' expr ')'
          = { call action2 }
```

In many applications, output is not done directly by the actions; rather, a data structure, such as a parse tree, is constructed in memory, and transformations are applied to it before output is generated. Parse trees are particularly easy to construct, given routines which build and maintain the tree structure desired. For example, suppose we have a C function "node", written so that the call

```
node( L, n1, n2 )
```

creates a node with label L, and descendants n1 and n2, and returns a pointer to the newly created node. Then we can cause a parse tree to be built by supplying actions such as:

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

in our specification.

The user may define other variables to be used by the actions. Declarations and definitions can appear in two places in the Yacc specification: in the declarations section, and at the head of the rules sections, before the first grammar rule. In each case, the declarations and definitions are enclosed in the marks "%{" and "%}". Declarations and definitions placed in the declarations section have global scope, and are thus known to the action statements and the lexical analyzer. Declarations and definitions placed at the head of the rules section have scope local to the action statements. Thus, in the above example, we might have included

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

in the declarations section, or, perhaps,

```
%{ static int variable; %}
```

at the head of the rules section. If we were writing Ratfor actions, we might want to include some COMMON statements at the beginning of the rules section, to allow for easy communication between the actions and other routines. For both C and Ratfor, Yacc has used only external names beginning in "yy"; the user should avoid such names.

Section 3: Lexical Analysis

The user must supply a lexical analyzer which reads the input stream and communicates tokens (with values, if desired) to the parser. The lexical analyzer is an integer valued function called yylex, in both C and Ratfor. The function returns an integer which represents the type of the token. The value to be associated in the parser with that token is assigned to the integer variable yylval. Thus, a lexical analyzer written in C should begin

```
yylex ( ) {
    extern int yylval;
    ...
```

while a lexical analyzer written in Ratfor should begin

```
integer function yylex(yylval)
integer yylval
...
```

Clearly, the parser and the lexical analyzer must agree on the type numbers in order for communication between them to take place. These numbers may be chosen by Yacc, or chosen by the user. In either case, the "define" mechanisms of C and Ratfor are used to allow the lexical analyzer to return these numbers symbolically. For example, suppose that the token name DIGIT has been defined in the declarations section of the specification. The relevant portion of the lexical analyzer (in C) might look like:

```
yylex() {
extern int yylval;
int c;
...
c = getchar( );
...
if( c >= '0' && c <= '9' ) {
    yylval = c-'0';
    return(DIGIT);
}
...
}
```

The relevant portion of the Ratfor lexical analyzer might look like:

```
integer function yylex(yylval)
integer yylval, digits(10), c
...
data digits(1) / "0" /;
data digits(2) / "1" /;
...
data digits(10) / "9" /;
...
# set c to the next input character
...
do i = 1, 10 {
    if(c .EQ. digits(i)) {
        yylval = i-1
        yylex = DIGIT
        return
    }
}
...
}
```

In both cases, the intent is to return a token type of DIGIT, and a value equal to the numerical value of the digit. Provided that the lexical analyzer code is placed in the programs section of the specification, the identifier DIGIT will be redefined to be equal to the type number associated with the token name DIGIT.

This mechanism leads to clear and easily modified lexical analyzers; the only pitfall is that it makes it important to avoid using any names in the grammar which are reserved or significant in the chosen language; thus, in both C and Ratfor, the use of token names of "if" or "yylex" will almost certainly cause severe difficulties when the lexical analyzer is compiled. The token name "error" is reserved for error handling, and should not be used naively (See Section 5).

As mentioned above, the type numbers may be chosen by Yacc or by the user. In the default situation, the numbers are chosen by Yacc. The default type number for a literal character is the numerical value of the character, considered as a 1 byte integer. Other token names are assigned type numbers starting at 257. It is a difficult, machine dependent operation to determine the numerical value of an input character in Ratfor (or Fortran). Thus, the Ratfor user of Yacc will probably wish to set his own type numbers, or not use any literals in his specification.

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

There is one exception to this situation. For sticky historical reasons, the endmarker must have type number 0. Note that this is not unattractive in C, since the nul character is returned upon end of file; in Ratfor, it makes no sense. This type number cannot be redefined by the user; thus, all lexical analyzers should be prepared to return 0 as a type number upon reaching the end of their input.

Section 4: Ambiguity, Conflicts, and Precedence

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

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

is a natural way of expressing the fact that one way of forming an arithmetic expression is to put two other expressions together with a minus sign between them. Unfortunately, this grammar rule does not completely specify the way that all complex inputs should be structured. For example, if we have input of the form

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

the rule would permit us to treat this input either as

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

or as

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

(We speak of the first as *left association* of operators, and the second as *right association*).

Yacc detects such ambiguities when it is attempting to build the parser. It is instructive to consider the problem that confronts the parser when it is given an input such as

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

When the parser has read the second *expr*, the input which it has seen.

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

matches the right side of the grammar rule above. One valid thing for the parser to do is to *reduce* the input it has seen by applying this rule; after applying the rule, it would have reduced the input it had already seen to *expr* (the left side of the rule). It could then read the final part of the input:

$$- \text{expr}$$

and again reduce by the rule. We see that the effect of this is to take the left associative interpretation.

Alternatively, when the parser has seen

expr - expr

it could defer the immediate application of the rule, and continue reading (the technical term is *shifting*) the input until it had seen

expr - expr - expr

It could then apply the grammar rule to the rightmost three symbols, reducing them to expr and leaving

expr - expr

Now it can reduce by the rule again; the effect is to take the right associative interpretation. Thus, having read

expr - expr

the parser can do two legal things, a shift or a reduction, and has no way of deciding between them. We refer to this as a *shift/reduce conflict*. It may also happen that the parser has a choice of two legal reductions; this is called a *reduce/reduce conflict*.

When there are shift/reduce or reduce/reduce conflicts, Yacc still produces a parser. It does this by selecting one of the valid steps wherever it has a choice. A rule which describes which choice to make in a given situation is called a *disambiguating rule*.

Yacc has two disambiguating rules which are invoked by default, in the absence of any user directives to the contrary:

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

Rule 1 implies that reductions are deferred whenever there is a choice, in favor of shifts. Rule 2 gives the user rather crude control over the behavior of the parser in this situation, but the proper use of reduce/reduce conflicts is still a black art, and is properly considered an advanced topic.

Conflicts may arise because of mistakes in input or logic, or because the grammar rules, while consistent, require a more complex parser than Yacc can construct. In these cases, the application of disambiguating rules is inappropriate, and leads to a parser which is in error. For this reason, Yacc always reports the number of shift/reduce and reduce/reduce conflicts which were resolved by Rule 1 and Rule 2.

In general, whenever it is possible to apply disambiguating rules to produce a correct parser, it is also possible to rewrite the grammar rules so that the same inputs are read, but there are no conflicts. For this reason, most previous systems like Yacc have considered conflicts to be fatal errors. Our experience has suggested that this rewriting is somewhat unnatural to do, and produces slower parsers; thus, Yacc will produce parsers even in the presence of conflicts.

As an example of the power of disambiguating rules, consider a fragment from a programming language involving an "if-then-else" construction:

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

Here, we consider IF and ELSE to be tokens, cond to be a nonterminal symbol describing conditional (logical) expressions, and stat to be a nonterminal symbol describing statements. In the following, we shall refer to these two rules as the *simple-if* rule and the *if-else* rule, respectively.

These two rules form an ambiguous construction, since input of the form

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

can be structured according to these rules in two ways:

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

or

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

The second interpretation is the one given in most programming languages which have this construct. Each ELSE is associated with the last preceding "un-ELSE'd" IF. In this example, consider the situation where the parser has seen

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

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

```
IF ( C1 ) stat
```

and then read the remaining input,

```
ELSE S2
```

and reduce

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

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

On the other hand, we may *shift* the ELSE and read S2, and then reduce the right hand portion of

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

by the if-else rule to get

```
IF ( C1 ) stat
```

which can be reduced by the simple-if rule. This leads to the second of the above groupings of the input, which is usually desired.

Once again the parser can do two valid things – we have a shift/reduce conflict. The application of disambiguating rule 1 tells the parser to shift in this case, which leads to the desired grouping.

Notice that this shift/reduce conflict arises only when there is a particular current input symbol, ELSE, and particular inputs already seen, such as

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

In general, there may be many conflicts, and each one will be associated with an input symbol and a set of previously read inputs. The previously read inputs are characterized by the *state* of the parser, which is assigned a nonnegative integer. The number of states in the parser is typically two to five times the number of grammar rules.

When Yacc is invoked with the verbose (-v) option (see Appendix B), it produces a file of user output which includes a description of the states in the parser. For example, the output corresponding to the above example might be:

23: shift/reduce Conflict (Shift 45, Reduce 18) on ELSE

State 23

```
stat : IF ( cond ) stat_  
stat : IF ( e_ ) stat_ELSE stat  
  
ELSE    shift 45  
        reduce 18
```

The first line describes the conflict, giving the state and the input symbol. The state title follows, and a brief description of the grammar rules which are active in this state. The underline “_” describes the portions of the grammar rules which have been seen. Thus in the example, in state 23 we have seen input which corresponds to

```
IF ( cond ) stat
```

and the two grammar rules shown are active at this time. The actions possible are, if the input symbol is ELSE, we may shift into state 45. In this state, we should find as part of the description a line of the form

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

because in this state we will have read and shifted the ELSE. Back in state 23, the alternative action, described by “_”, is to be done if the input symbol is not mentioned explicitly in the above actions; thus, in this case, if the input symbol is not ELSE, we should reduce by grammar rule 18, which is presumably

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

Notice that the numbers following “shift” commands refer to other states, while the numbers following “reduce” commands refer to grammar rule numbers. In most states, there will be only one reduce action possible in the state, and this will always be the default command. The user who encounters unexpected shift/reduce conflicts will probably want to look at the verbose output to decide whether the default actions are appropriate. In really tough cases, the user might need to know more about the behavior and construction of the parser than can be covered here; in this case, a reference such as [1] might be consulted; the services of a local guru might also be appropriate.

There is one common situation where the rules given above for resolving conflicts are not sufficient; this is in the area of arithmetic expressions. Most of the commonly used constructions for arithmetic expressions can be naturally described by the notion of *precedence* levels for operators, together with information about left or right associativity. It turns out that ambiguous grammars with appropriate disambiguating rules can be used to create parsers which are faster and easier to write than parsers constructed from unambiguous grammars. The basic notion is to write grammar rules of the form

```
expr : expr OP expr
```

and

```
expr : UNARY expr
```

for all binary and unary operators desired. This creates a very ambiguous grammar, with many

parsing conflicts. As disambiguating rules, the user specifies the precedence, or binding strength, of all the operators, and the associativity of the binary operators. This information is sufficient to allow Yacc to resolve the parsing conflicts in accordance with these rules, and construct a parser which realizes the desired precedences and associativities.

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

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

describes the precedence and associativity of the four arithmetic operators. Plus and minus are left associative, and have lower precedence than star and slash, which are also left associative. The keyword %right is used to describe right associative operators, and the keyword %nonassoc is used to describe operators, like the operator .LT. in Fortran, which may not associate with themselves; thus,

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

is illegal in Fortran, and such an operator would be described with the keyword %nonassoc in Yacc. As an example of the behavior of these declarations, the description

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

```
%%
```

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

might be used to structure the input

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

as follows:

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

When this mechanism is used, unary operators must, in general, be given a precedence. An interesting situation arises when we have a unary operator and a binary operator which have the same symbolic representation, but different precedences. An example is unary and binary '-'; frequently, unary minus is given the same strength as multiplication, or even higher, while binary minus has a lower strength than multiplication. We can indicate this situation by use of another keyword, %prec, to change the precedence level associated with a particular grammar rule. %prec appears immediately after the body of the grammar rule, before the action or closing semicolon, and is followed by a token name or literal; it causes the precedence of the grammar rule to become that of the token name or literal. Thus, to make unary minus have the same precedence as multiplication, we might write:

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

Notice that the precedences which are described by %left, %right, and %nonassoc are independent of the declarations of token names by %token. A symbol can be declared by %token, and, later in the declarations section, be given a precedence and associativity by one of the above methods. It is true, however, that names which are given a precedence or associativity are also declared to be token names, and so in general do not need to be declared by %token, although it does not hurt to do so.

As we mentioned above, the precedences and associativities are used by Yacc to resolve parsing conflicts; they give rise to disambiguating rules. Formally, the rules work as follows:

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

There are a number of points worth making about this use of disambiguation. There is no reporting of conflicts which are resolved by this mechanism, and these conflicts are not counted in the number of shift/reduce and reduce/reduce conflicts found in the grammar. This means that occasionally mistakes in the specification of precedences disguise errors in the input grammar; it is a good idea to be sparing with precedences, and use them in an essentially "cookbook" fashion, until some experience has been gained. Frequently, not enough operators or precedences have been specified; this leads to a number of messages about shift/reduce or reduce/reduce conflicts. The cure is usually to specify more precedences, or use the %prec mechanism, or both. It is generally good to examine the verbose output file to ensure that the conflicts which are being reported can be validly resolved by precedence.

Section 5: Error Handling

Error handling is an extremely difficult area, and many of the problems are semantic ones. When an error is found, for example, it may be necessary to reclaim parse tree storage, delete or alter symbol table entries, and, typically, set switches to avoid putting out any further output.

It is generally not acceptable to stop all processing when an error is found; we wish to continue scanning the input to find any further syntax errors. This leads to the problem of getting the parser "restarted" after an error. The general class of algorithms to do this involves reading ahead and discarding a number of tokens from the input string, and attempting to adjust the parser so that input can continue.

To allow the user some control over this process, Yacc provides a simple, but reasonably general, feature. The token name "error" is reserved for error handling. This name can be used in grammar rules; in effect, it suggests places where errors are expected, and recovery might take place. The parser attempts to find the last time in the input when the special token "error" is permitted. The parser then behaves as though it saw the token name "error" as an input token, and attempts to parse according to the rule encountered. The token at which the error was detected remains the next input token after this error token is processed. If no special error rules have been specified, the processing effectively halts when an error is detected.

In order to prevent a cascade of error messages, the parser assumes that, after detecting an error, it remains in error state until three tokens have been successfully read and shifted. If an error is detected when the parser is already in error state, no error message is given, and the input token is quietly deleted.

As a common example, the user might include a rule of the form

```
statement : error ;
```

in his specification. This would, in effect, mean that on a syntax error the parser would attempt to skip over the statement in which the error was seen. (Notice, however, that it may be difficult or impossible to tell the end of a statement, depending on the other grammar rules). More precisely, the parser will scan ahead, looking for three tokens that might legally follow a statement, and start processing at the first of these; if the beginnings of statements are not sufficiently distinctive, it may make a false start in the middle of a statement, and end up reporting a second error where there is in fact no error.

The user may supply actions after these special grammar rules, just as after the other grammar rules. These actions might attempt to reinitialize tables, reclaim symbol table space, etc.

The above form of grammar rule is very general, but somewhat difficult to control. Somewhat easier to deal with are rules of the form

```
statement : error ';' ;
```

Here, when there is an error, the parser will again attempt to skip over the statement, but in this case will do so by skipping to the next ";". All tokens after the error and before the next ";" give syntax errors, and are discarded. When the ";" is seen, this rule will be reduced, and any "cleanup" action associated with it will be performed.

Still another form of error rule arises in interactive applications, where we may wish to prompt the user who has incorrectly input a line, and allow him to reenter the line. In C we might write:

```
inputline: error '\n' prompt inputline
          = { $$ = $4; };
```

```
prompt: /* matches no input */
        = { printf( "Reenter last line: " ); };
```

There is one difficulty with this approach; the parser must correctly process three input tokens before it is prepared to admit that it has correctly resynchronized after the error. Thus, if the reentered line contains errors in the first two tokens, the parser will simply delete the offending tokens, and give no message; this is clearly unacceptable. For this reason, there is a mechanism in both C and Ratfor which can be used to force the parser to believe that resynchronization has taken place. One need only include a statement of the form

```
yyerrok ;
```

in his action after such a grammar rule, and the desired effect will take place; this name will be expanded, using the "# define" mechanism of C or the "define" mechanism of Ratfor, into an appropriate code sequence. For example, in the situation discussed above where we want to prompt the user to produce input, we probably want to consider that the original error has been recovered when we have thrown away the previous line, including the newline. In this case, we can reset the error state before putting out the prompt message. The grammar rule for the nonterminal symbol prompt becomes:

```
prompt: /* matches no input */
        = {
            yyerrok;
            printf( "Reenter last line: " );
        };
```

There is another special feature which the user may wish to use in error recovery. As mentioned above, the token seen immediately after the "error" symbol is the input token at which the error was discovered. Sometimes, this is seen to be inappropriate; for example, an error recovery action might take upon itself the job of finding the correct place to resume input. In this case, the user wishes a way of clearing the previous input token held in the parser. One need only include a statement of the form

```
yyclearin ;
```

in his action; again, this expands, in both C and Ratfor, to the appropriate code sequence. For example, suppose the action after error were to call some sophisticated resynchronization routine, supplied by the user, which attempted to advance the input to the beginning of the next valid statement. After this routine was called, the next token returned by yylex would presumably be the first token in a legal statement; we wish to throw away the old, illegal token, and reset the error state. We might do this by the sequence:

```
statement : error
          = {
            resynch( );
            yyerrok ;
            yyclearin ;
          };
```

These mechanisms are admittedly crude, but do allow for a simple, fairly effective recovery of the parser from many errors, and have the virtue that the user can get "handles" by which he can deal with the error actions required by the lexical and output portions of the system.

Section 6C: The C Language Yacc Environment

The default mode of operation in Yacc is to write actions and the lexical analyzer in C. This has a number of advantages; primarily, it is easier to write character handling routines, such as the lexical analyzer, in a language which supports character-by-character I/O, and has shifting and masking operators.

When the user inputs a specification to Yacc, the output is a file of C programs, called "y.tab.c". These are then compiled, and loaded with a library; the library has default versions of a number of useful routines. This section discusses these routines, and how the user can write his own routines if desired. The name of the Yacc library is system dependent; see Appendix B.

The subroutine produced by Yacc is called "yyparse"; it is an integer valued function. When it is called, it in turn repeatedly calls "yylex", the lexical analyzer supplied by the user (see Section 3), to obtain input tokens. Eventually, either an error is detected, in which case (if no error recovery is possible) yyparse returns the value 1, or the lexical analyzer returns the endmarker token (type number 0), and the parser accepts. In this case, yyparse returns the value 0.

Three of the routines on the Yacc library are concerned with the "external" environment of yyparse. There is a default "main" program, a default "initialization" routine, and a default "accept" routine, respectively. They are so simple that they will be given here in their entirety:

```
main( argc, argv )
int argc;
char *argv[ ]
{
    yyinit( argc, argv );
    if( yyparse( ) )
        return;
    yyacct( );
}

yyinit( ) { }

yyacct( ) { }
```

By supplying his own versions of yyinit and/or yyacct, the user can get control either before the parser is called (to set options, open input files, etc.) or after the accept action has been done (to close files, call the next pass of the compiler, etc.). Note that yyinit is called with the two "command line" arguments which have been passed into the main program. If neither of these routines is redefined, the default situation simply looks like a call to the parser, followed by the termination of the program. Of course, in many cases the user will wish to supply his own main program; for example, this is necessary if the parser is to be called more than once.

The other major routine on the library is called "yyerror"; its main purpose is to write out a message when a syntax error is detected. It has a number of hooks and handles which attempt to make this error message general and easy to understand. This routine is somewhat more complex, but still approachable:

```
extern int yyline, /* input line number */

yyerror(s)
char *s;
{
    extern int yychar;
    extern char *yyterm[ ];

    printf("\n%s", s );
    if( yyline )
        printf(" line %d,", yyline );
    printf(" on input: ");
    if( yychar >= 0400 )
        printf("%s\n", yyterm[yychar-0400] );
    else switch ( yychar ) {
    case `t`: printf( "\\t\n" ); return;
    case `n`: printf( "\\n\n" ); return;
    case `0`: printf( "$end\n" ); return;
    default: printf( "%c\n", yychar ); return;
    }
}
```

The argument to `yyerror` is a string containing an error message; most usually, it is "syntax error". `yyerror` also uses the external variables `yyline`, `yychar`, and `yyterm`. `yyline` is a line number which, if set by the user to a nonzero number, will be printed out as part of the error message. `yychar` is a variable which contains the type number of the current token. `yyterm` has the names, supplied by the user, for all the tokens which have names. Thus, the routine spends most of its time trying to print out a reasonable name for the input token. The biggest problem with the routine as given is that, on Unix, the error message does not go out on the error file (file 2). This is hard to arrange in such a way that it works with both the portable I/O library and the system I/O library; if a way can be worked out, the routine will be changed to do this. *Beware*: This routine will not work if any token names have been given redefined type numbers. In this case, the user must supply his own `yyerror` routine. Hopefully, this "feature" will disappear soon.

Finally, there is another feature which the C user of Yacc might wish to use. The integer variable `yydebug` is normally set to 0. If it is set to 1, the parser will output a verbose description of its actions, including a discussion of which input symbols have been read, and what the parser actions are. Depending on the operating environment, it may be possible to set this variable by using a debugging system.

Section 6R: The Ratfor Language Yacc Environment

For reasons of portability or compatibility with existing software, it may be desired to use Yacc to generate parsers in Ratfor, or, by extension, in portable Fortran. The user is likely to work considerably harder doing this than he might if he were to use C.

When the user inputs a specification to Yacc, and specifies the Ratfor option (see Appendix B), the output is a file of Ratfor programs called "y.tab.r". These programs are then compiled, and provide the desired subroutine.

The subroutine produced by Yacc which does the input process is an integer function called "yypars". When it is called, it in turn repeatedly calls "yylex", the lexical analyzer supplied by the user (see Section 3). Eventually, either an error is detected, in which case (if no error recovery is possible) `yypars` returns the value 1, or the lexical analyzer returns the end-marker (type number 0), and the parser accepts. In this case, `yypars` returns 0.

Unlike the C program situation (see Section 6C) there is no library of Ratfor routines which must be used in the loading process. As a side effect of this, *the user must supply a main program which calls yypars*. A suggested Ratfor main program is

```
integer yypars
n = yypars(0)
if( n .EQ. 0 ) {
    ... here if the program accepted
} else {
    ... here if there were unrecoverable errors
}
end
```

Notice that there is no easy way for the user to get control when an error is detected, since the Fortran language provides only a very crude character string capability.

There is another feature which the Ratfor user might wish to use. The argument to yypars is normally 0. If it is set to 1, the parser will output a verbose description of its actions, including a discussion of which input symbols have been read, and what the parser actions are. During the input process, the value of this debug flag is kept in a common variable yydebu, which is available to the actions and may be set and reset at will.

Statement labels 1 through 1000 are reserved for the parser, and may not appear in actions; note that, because Ratfor has a more modern control structure than Fortran, it is rarely necessary to use statement labels at all; the most frequent use of labels in Ratfor is in formatted I/O.

Because Fortran has no standard character set and not even a standard character width, it is difficult to produce a lexical analyzer in portable Fortran. The usual solution is to provide a routine which does a table search to get the internal type number for each input character, with the understanding that such a routine can be recoded to run far faster for any particular machine.

Finally, we must warn the user that the Ratfor feature of Yacc has been operational for a much shorter time than the other portions of the system. If past experience is any guide, the Ratfor support will develop and become more powerful and better human engineered in response to user complaints and requirements. Thus, the potential Ratfor user might do well to contact the author to discuss his own particular needs.

Section 7. Hints for Preparing Specifications

This section contains miscellaneous hints on preparing efficient, easy to change, and clear specifications. The individual subsections are, more or less, independent; the reader seeing Yacc for the first time may well find that this entire section could be omitted.

Input Style

It is difficult to input rules with substantial actions and still have a readable specification file. The following style hints owe much to Brian Kernighan, and are officially endorsed by the author.

- a. Use all capital letters for token names, all lower case letters for nonterminal names. This rule comes under the heading of "knowing who to blame when things go wrong."
- b. Put grammar rules and actions on separate lines. This allows either to be changed without an automatic need to change the other.

- c. Put all rules with the same left hand side together. Put the left hand side in only once, and let all following rules begin with a vertical bar.
- d. Indent rule bodies by one tab stop, and action bodies by two tab stops.

The example in Appendix A is written following this style, as are the examples in the text of this paper (where space permits). The user must make up his own mind about these stylistic questions; the central problem, however, is to make the rules visible through the morass of action code.

Common Actions

When several grammar rules have the same action, the user might well wish to provide only one code sequence. A simple, general mechanism is, of course, to use subroutine calls. It is also possible to put a label on the first statement of an action, and let other actions be simply a goto to this label. Thus, if the user had a routine which built trees, he might wish to have only one call to it, as follows:

```
expr :
  expr '+' expr =
  { binary:
    $$ = btree( $1, $2, $3 );
  }
  expr '-' expr =
  {
    goto binary;
  }
  expr '*' expr =
  {
    goto binary;
  }
  ;
```

Left Recursion

The algorithm used by the Yacc parser encourages so called "left recursive" grammar rules: rules of the form

```
name : name rest_of_rule ;
```

These rules frequently arise when writing specifications of sequences and lists:

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

and

```
sequence :
  item |
  sequence item ;
```

Notice that, in each of these cases, the first rule will be reduced for the first item only, and the second rule will be reduced for the second and all succeeding items.

If the user were to write these rules right recursively, such as

```
sequence :
  item |
  item sequence ;
```

the parser would be a bit bigger, and the items would be seen, and reduced, from right to left. More seriously, an internal stack in the parser would be in danger of overflowing if a very long sequence were read. Thus, the user should use left recursion wherever reasonable.

The user should also consider whether a sequence with zero elements has any meaning, and if so, consider writing the sequence specification with an empty rule:

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

Once again, the first rule would always be reduced exactly once, before the first item was read, and then the second rule would be reduced once for each item read. Experience suggests that permitting empty sequences leads to increased generality, which frequently is not evident at the time the rule is first written. There are cases, however, when the Yacc algorithm can fail when such a change is made. In effect, conflicts might arise when Yacc is asked to decide which empty sequence it has seen, when it hasn't seen enough to know! Nevertheless, this principle is still worth following wherever possible.

Lexical Tie-ins

Frequently, there are lexical decisions which depend on the presence of various constructions in the specification. For example, the lexical analyzer might want to delete blanks normally, but not within quoted strings. Or names might be entered into a symbol table in declarations, but not in expressions.

One way of handling these situations is to create a global flag which is examined by the lexical analyzer, and set by actions. For example, consider a situation where we have a program which consists of 0 or more declarations, followed by 0 or more statements. We declare a flag called "dflag", which is 1 during declarations, and 0 during statements. We may do this as follows:

```
%{  
    int dflag ;  
}%  
%%  
program :  
    decls stats ;  
  
decls :  
    = /* empty */  
    {  
        dflag = 1 ;  
    }  
    decls declaration ;  
  
stats :  
    = /* empty */  
    {  
        dflag = 0 ;  
    }  
    stats statement ;  
  
... other rules ...
```

The flag dflag is now set to zero when reading statements, and 1 when reading declarations, *except for the first token in the first statement*. This token must be seen by the parser before it can

tell that the declaration section has ended and the statements have begun. Frequently, however, this single token exception does not affect the lexical scan required.

Clearly, this kind of "backdoor" approach can be elaborated on to a noxious degree. Nevertheless, it represents a way of doing some things that are difficult, if not impossible, to do otherwise.

Bundling

Bundling is a technique for collecting together various character strings so that they can be output at some later time. It is derived from a feature of the same name in the compiler/compiler TMG [6].

Bundling has two components – a nice user interface, and a clever implementation trick. They will be discussed in that order.

The user interface consists of two routines, "bundle" and "bprint".

```
bundle( a1, a2, . . . , an )
```

accepts a variable number of arguments which are either character strings or bundles, and returns a bundle, whose value will be the concatenation of the values of a1, . . . , an.

```
bprint( b )
```

accepts a bundle as argument and outputs its value.

For example, suppose that we wish to read arithmetic expressions, and output function calls to routines called "add", "sub", "mul", "div", and "assign". Thus, we wish to translate

$$a = b - c*d$$

into

```
assign(a,sub(b,mul(c,d)))
```

A Yacc specification file which does this is given in Appendix D; this includes an implementation of the bundle and bprint routines. A rule and action of the form

```
expr:
  expr '+' expr =
  {
    $$ = bundle( "add(", $1, ",", $3, ")" );
  }
```

causes the returned value of expr to be come a bundle, whose value is the character string containing the desired function call. Each NAME token has a value which is a pointer to the actual name which has been read. Finally, when the entire input line has been read and the value has been bundled, the value is written out and the bundles and names are cleared, in preparation for the next input line.

Bundles are implemented as arrays of pointers, terminated by a zero pointer. Each pointer either points to a bundle or to a character string. There is an array, called *bundle space*, which contains all the bundles.

The implementation trick is to check the values of the pointers in bundles – if the pointer points into bundle space, it is assumed to point to a bundle; otherwise it is assumed to point to a character string.

The treatment of functions with a variable number of arguments, like bundle, is likely to differ from one implementation of C to another.

In general, one may wish to have a simple storage allocator which allocates and frees bundles, in order to handle situations where it is not appropriate to completely clear all of bundle space at one time.

Reserved Words

Some programming languages permit the user to use words like "if", which are normally reserved, as label or variable names, provided that such use does not conflict with the legal use of these names in the programming language. This is extremely hard to do in the framework of Yacc, since it is difficult to pass the required information to the lexical analyzer which tells it "this instance of if is a keyword, and that instance is a variable". The user can make a stab at it, using the mechanism described in the last subsection, but it is difficult.

A number of ways of making this easier are under advisement, and one will probably be supported eventually. Until this day comes, I suggest that the keywords be *reserved*; that is, be forbidden for use as variable names. There are powerful stylistic reasons for preferring this, anyway (he said weakly . . .).

Non-integer Values

Frequently, the user wishes to have values which are bigger than integers; again, this is an area where Yacc does not make the job as easy as it might, and some additional support is likely. Nevertheless, at the cost of writing a storage manager, the user can return pointers or indices to blocks of storage big enough to contain the full values desired.

Previous Work

There have been many previous applications of Yacc. The user who is contemplating a big application might well find that others have developed relevant techniques, or even portions of grammars. Yacc specifications appear to be easier to change than the equivalent computer programs, so that the "prior art" is more relevant here as well.

Section 8: User Experience, Summary, and Acknowledgements

Yacc has been used in the construction of a C compiler for the Honeywell 6000, a system for typesetting mathematical equations, a low level implementation language for the PDP 11, APL and Basic compilers to run under the UNIX system, and a number of other applications.

To summarize, Yacc can be used to construct parsers; these parsers can interact in a fairly flexible way with the lexical analysis and output phases of a larger system. The system also provides an indication of ambiguities in the specification, and allows disambiguating rules to be supplied to resolve these ambiguities.

Because the output of Yacc is largely tables, the system is relatively language independent. In the presence of reasonable applications, Yacc could be modified or adapted to produce subroutines for other machines and languages. In addition, we continue to seek better algorithms to improve the lexical analysis and code generation phases of compilers produced using Yacc.

This document would be incomplete if I did not give credit to a most stimulating collection of users, who have goaded me beyond my inclination, and frequently beyond my ability, in their endless search for "one more feature". Their irritating unwillingness to learn how to do things my way has usually led to my doing things their way; most of the time, they have been right. B. W. Kernighan, P. J. Plauger, S. I. Feldman, C. Imagna, M. E. Lesk, and A. Snyder will recognize some of their ideas in the current version of Yacc. Al Aho also deserves recognition for bringing the mountain to Mohammed, and other favors.

References

- 1 Aho, A.V. and Johnson, S.C., "LR Parsing", Computing Surveys, Vol 6, No 2, June 1974, pp. 99-124.
- 2 Aho, A.V., Johnson, S.C., and Ullman, J.D., "Deterministic Parsing of Ambiguous Grammars", Proceedings of the A.C.M. Symposium on Principles of Programming Languages, October 1973, pp. 1-21; to appear in CACM.
- 3 Aho, A.V. and Ullman, J.D., Theory of Parsing, Translation, and Compiling. Volume 1 (1972) and Volume 2 (1973), Prentice-Hall, Englewood Cliffs, N.J.
- 4 Kernighan, B. W., Ratfor, a Rational Fortran
- 5 Ryder, B. B., "The PFORT Verifier," Software—Practice and Experience, Vol 4 (1974), pp 359-377.
- 6 McIlroy, M. D., A Manual for the TMG Compiler-writing Language
- 7 Ritchie, D. M., C Reference Manual

Appendix A: A Simple Example

This example gives the complete Yacc specification for a small desk calculator; the desk calculator has 26 registers, labeled a through z, and accepts arithmetic expressions made up of the operators +, -, *, /, % (mod operator), & (bitwise and), | (bitwise or), and assignment. If an expression is an assignment at the top level, the value is not printed; otherwise it is. As in C, an integer which begins with 0 (zero) is assumed to be octal; otherwise, it is assumed to be decimal.

As an example of a Yacc specification, the desk calculator does a reasonable job of showing the way that precedences and ambiguities are used, as well as showing how simple error recovery operates. The major oversimplifications are that the lexical analysis phase is much simpler than for most applications, and the output is produced immediately, line by line. Note the way that decimal and octal integers are read in by the grammar rules; frequently, this job is better done by the lexical analyzer.

```
%token DIGIT LETTER /* these are token names */
%left ↑ /* declarations of operator precedences */
%left '&'
%left '+' '-'
%left '*' '/' '%'
%left UMINUS /* supplies precedence for unary minus */
%{ /* declarations used by the actions */
    int base;
    int regs[26];
%}

%% /* beginning of rules section */

list : /* list is the start symbol */
    | /* empty */
    list stat '\n' |
    list error '\n' =
    {
        yyerrok ;
    } ;

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

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

```
expr '+' expr =
{
    $$ = $1 + $3 ;
}
expr '-' expr =
{
    $$ = $1 - $3 ;
}
expr '*' expr =
{
    $$ = $1 * $3 ;
}
expr '/' expr =
{
    $$ = $1 / $3 ;
}
expr '%' expr =
{
    $$ = $1 % $3 ;
}
expr '&' expr
{
    $$ = $1 & $3 ;
}
expr '|' expr
{
    $$ = $1 | $3 ;
}
'-' expr %prec UMINUS
{
    $$ = - $2 ;
}
LETTER
{
    $$ = regs[$1] ;
}
number ;
```

```
number :
DIGIT =
{
    $$ = $1 ;
    base = 10 ;
    if( $1 == 0 )
        base = 8 ;
}
number DIGIT =
{
    $$ = base * $1 + $2 ;
};
```

```
%% /* start of programs */
```

```
yylex() /* lexical analysis routine */
{
    /* returns LETTER for a lower case letter, yylval = 0 through 25 */
    /* return DIGIT for a digit, yylval = 0 through 9 */
    /* all other characters are returned immediately */

    int c ;

    while( (c=getchar( )) != ' ' )
        ;
    if( c >= 'a' && c <= 'z' ) {
        yylval = c - 'a' ;
        return( LETTER ) ;
    }
    if( c >= '0' && c <= '9' ) {
        yylval = c - '0' ;
        return( DIGIT ) ;
    }
    return( c ) ;
}
```


Appendix B: Use of Yacc on Unix

Suppose that the Yacc specification is on a file called yfile. If the actions are in C, Yacc is invoked by

```
yacc yfile
```

The output appears on file y.tab.c. To compile the parser and load it with the Yacc library, use the command

```
cc y.tab.c -ly
```

If Yacc is invoked with the option `-v`:

```
yacc -v yfile
```

a verbose description of the parser is produced on file y.output. The C user should consult section 6C for more information about the run time environment.

If the actions are in Ratfor, the user should invoke Yacc with the option `-r`:

```
yacc -r yfile
```

The Ratfor output appears on file y.tab.r. It may be compiled by

```
rc -2 y.tab.r
```

Note that when Yacc is used to produce Ratfor programs, there is no need to load these programs with any library.

If the `-v` action is also invoked:

```
yacc -rv yfile
```

a verbose description of the parser is produced on file y.output. The Ratfor user should consult section 6R for more information about the run time environment.

Appendix C: Old Features Supported but not Encouraged

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

1. Literals may be delimited by double quotes "" as well as single quotes ''.
2. Literals may be more than one character long. If all the characters are alphabetic, numeric, or `_`, the type number of the literal is defined, just as if the literal did not have the quotes around it. Otherwise, it is difficult to find the value for such literals.

The use of multi-character literals is likely to mislead those unfamiliar with Yacc, since it suggests that Yacc is doing a job which must be actually done by the lexical analyzer.

3. Most places where `%` is legal, backslash "\ " may be used. In particular, `\\` is the same as `\"`, `%left` the same as `%left`, etc.
4. There are a number of other synonyms:
 - `%<` is the same as `%left`
 - `%>` is the same as `%right`
 - `%binary` and `%2` are the same as `%nonassoc`
 - `%0` and `%term` are the same as `%token`
 - `%=` is the same as `%prec`
5. The curly braces "{" and "}" around an action are optional if the action consists of a single C statement. (They are always required in Ratfor).

Appendix D: An Example of Bundling

The following program is an example of the technique of bundling; this example is discussed in Section 7.

```
/* warnings:
```

1. This works on Unix; the handling of functions with a variable number of arguments is different on different systems.
2. A number of checks for array bounds have been left out to avoid obscuring the basic ideas, but should be there in a practical program.

```
*/
```

```
%token NAME
```

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

```
%%
```

```
lines :
```

```
    = /* empty */  
    {  
        bclear() ;  
    }  
    |  
    lines expr '\n' =  
    {  
        bprint( $2 ) ;  
        printf( "\n" ) ;  
        bclear() ;  
    }  
    |  
    lines error '\n' =  
    {  
        bclear() ;  
        yyerrok ;  
    } ;
```

```
expr :
```

```
    expr '+' expr =  
    {  
        $$ = bundle( "add(", $1, ",", $3, ")" ) ;  
    }  
    |  
    expr '-' expr =  
    {  
        $$ = bundle( "sub(", $1, ",", $3, ")" ) ;  
    }  
    |  
    expr '*' expr =  
    {  
        $$ = bundle( "mul(", $1, ",", $3, ")" ) ;  
    }  
    |
```

```
    expr '/' expr,
    {
        $$ = bundle( "div(", $1, ",", $3, ")" );
    }
    '(' expr ')' =
    {
        $$ = $2;
    }
    NAME '=' expr =
        $$ = bundle( "assign(", $1, "=", $3, ")" );
    }
    NAME ;

%%

#define nsize 200
char names[nsize], *nptr { names };

#define bsize 500
int bspace[bsize], *bptr { bspace };

yylex()
{
    int c;

    c = getchar();
    while( c == ' ' )
        c = getchar();
    if( c >= 'a' && c <= 'z' ) {
        yylval = nptr;
        for( ; c >= 'a' && c <= 'z'; c=getchar() )
            *nptr++ = c;
        ungetc( c );
        *nptr++ = '\0';
        return( NAME );
    }
    return( c );
}

bclear()
{
    nptr = names;
    bptr = bspace;
}

bundle( a1,a2,a3,a4,a5 )
{
    int i, j, *p, *obp;

    p = &a1;
    i = nargs();
```

```
    obp = bptr;

    for( j=0; j<i; ++j )
        *bptr++ = *p++;
    *bptr++ = 0;
    return( obp );
}

bprint( p )
int *p;
{
    if( p >= bspace && p < &bspace[bsize] ) /* bundle */
        while( *p != 0 )
            bprint( *p++ );
    else printf( "%s", p );
}
```





Lex - A Lexical Analyzer Generator

M. E. Lesk and E. Schmidt

Bell Laboratories

Murray Hill, New Jersey 07974

Lex helps write programs whose control flow is directed by instances of regular expressions in the input stream. It is well suited for editor-script type transformations and for segmenting input in preparation for a parsing routine.

Lex source is a table of regular expressions and corresponding program fragments. The table is translated to a program which reads an input stream, copying it to an output stream and partitioning the input into strings which match the given expressions. As each such string is recognized the corresponding program fragment is executed. The recognition of the expressions is performed by a deterministic finite automaton generated by Lex. The program fragments written by the user are executed in the order in which the corresponding regular expressions occur in the input stream.

The lexical analysis programs written with Lex accept ambiguous specifications and choose the longest match possible at each input point. If necessary, substantial lookahead is performed on the input, but the input stream will be backed up to the end of the current partition, so that the user has general freedom to manipulate it.

Lex can be used to generate analyzers in either C or Ratfor, a language which can be translated automatically to portable Fortran. It is available on the PDP-11 UNIX, Honeywell GCOS, and IBM OS systems. Lex is designed to simplify interfacing with Yacc, for those with access to this compiler-compiler system.

Table of Contents

1. Introduction.	1
2. Lex Source.	3
3. Lex Regular Expressions.	3
4. Lex Actions.	5
5. Ambiguous Source Rules.	7
6. Lex Source Definitions.	8
7. Usage.	8
8. Lex and Yacc.	9
9. Examples.	10
10. Left Context Sensitivity.	11
11. Character Set.	12
12. Summary of Source Format.	12
13. Caveats and Bugs.	13
14. Acknowledgments.	13
15. References.	13

1 Introduction.

Lex is a program generator designed for lexical processing of character input streams. It accepts a high-level, problem oriented specification for character string matching, and produces a program in a general purpose language which recognizes regular expressions. The regular expressions are specified by the user in the source specifications given to Lex. The Lex written code recognizes these expressions in an input stream and partitions the input stream into strings matching the expressions. At the boundaries between strings program sections provided by the user are executed. The Lex source file asso-

ciates the regular expressions and the program fragments. As each expression appears in the input to the program written by Lex, the corresponding fragment is executed.

The user supplies the additional code beyond expression matching needed to complete his tasks, possibly including code written by other generators. The program that recognizes the expressions is generated in the general purpose programming language employed for the user's program fragments. Thus, a high level expression language is provided to write the string expressions to be matched while the user's freedom to write actions is unimpaired. This avoids forcing the user who wishes to use a string manipulation language for input analysis to

Source → [Lex] → yylex

Input → [yylex] → Output

An overview of Lex

Figure 1

write processing programs in the same and often inappropriate string handling language.

Lex is not a complete language, but rather a generator representing a new language feature which can be added to different programming languages, called "host languages." Just as general purpose languages can produce code to run on different computer hardware, Lex can write code in different host languages. The host language is used for the output code generated by Lex and also for the program fragments added by the user. Compatible run-time libraries for the different host languages are also provided. This makes Lex adaptable to different environments and different users. Each application may be directed to the combination of hardware and host language appropriate to the task, the user's background, and the properties of local implementations. At present there are only two host languages, C[1] and Fortran (in the form of the Ratfor language[2]). Lex itself exists on UNIX, GCOS, and OS/370; but the code generated by Lex may be taken anywhere the appropriate compilers exist.

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

For a trivial example, consider a program to delete from the input all blanks or tabs at the ends of lines.

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

is all that is required. The program contains a %% delimiter to mark the beginning of the rules, and one rule.


This rule contains a regular expression which matches one or more instances of the characters blank or tab (written \t for visibility, in accordance with the C language convention) just prior to the end of a line. The brackets indicate the character class made of blank and tab; the + indicates "one or more ..."; and the \$ indicates "end of line," as in QED. No action is specified, so the program generated by Lex (*yylex*) will ignore these characters. Everything else will be copied. To change any remaining string of blanks or tabs to a single blank, add another rule:

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

The finite automaton generated for this source will scan for both rules at once, observing at the termination of the string of blanks or tabs whether or not there is a newline character, and executing the desired rule action. The first rule matches all strings of blanks or tabs at the end of lines, and the second rule all remaining strings of blanks or tabs.

Lex can be used alone for simple transformations, or for analysis and statistics gathering on a lexical level. Lex can also be used with a parser generator to perform the lexical analysis phase; it is particularly easy to interface Lex and Yacc [3]. Lex programs recognize only regular expressions; Yacc writes parsers that accept a large class of context free grammars, but require a lower level analyzer to recognize input tokens. Thus, a combination of Lex and Yacc is often appropriate. When used as a preprocessor for a later parser generator, Lex is used to partition the input stream, and the parser generator assigns structure to the resulting pieces. The flow of control in such a case (which might be the first half of a compiler, for example) is shown in Figure 2. Additional programs, written by other generators or by hand, can be added easily to programs written by Lex. Yacc users will realize that the name *yylex* is what Yacc expects its lexical analyzer to be named; so that the use of this name by Lex simplifies interfacing.

Lex generates a deterministic finite automaton from the regular expressions in the source [4]. The automaton is interpreted, rather than compiled, in order to save space. The result is still a fast analyzer. In particular, the time



Lex with Yacc

Figure 2

taken by a Lex program to recognize and partition an input stream is proportional to the length of the input. The number of Lex rules or the complexity of the rules is not **important** in determining speed, unless rules which include forward context require a significant amount of re-scanning. What does increase with the number and complexity of rules is the size of the finite automaton, and therefore the size of the program generated by Lex.

In the program written by Lex, the user's fragments (representing the *actions* to be performed as each regular expression is found) are gathered as cases of a switch (in C) or branches of a computed GOTO (in Ratfor). The automaton interpreter directs the control flow. Opportunity is provided for the user to insert either declarations or additional statements in the routine containing the actions, or to add subroutines outside this action routine.

Lex is not limited to source which can be interpreted on the basis of one character lookahead. For example, if there are two rules, one looking for *ab* and another for *abcdefg*, and the input stream is *abcdefh*, Lex will recognize *ab* and leave the input pointer just before *cd*. . . Such backup is more costly than the processing of simpler languages.

2 Lex Source.

The general format of Lex source is:

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

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

```
%%
```

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

In the outline of Lex programs shown above, the *rules* represent the user's control decisions; they are a table, in which the left column contains *regular expressions* (see section 3) and the right column contains *actions*, program fragments to be executed when the expressions are recognized. Thus an individual rule might appear

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

to look for the string *integer* in the input stream and print the message "found keyword INT" whenever it appears. In this example the host procedural language is C and the C library function *printf* is used to print the string. The end of the expression is indicated by the first blank or tab character. If the action is merely a single C expression, it can just be given on the right side of the line; if it is compound, or takes more than a line, it should be enclosed in

braces. As a slightly more useful example, suppose it is desired to change a number of words from British to American spelling. Lex rules such as

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

would be a start. These rules are not quite enough, since the word *petroleum* would become *gaseum*; a way of dealing with this will be described later.

3 Lex Regular Expressions.

The definitions of regular expressions are very similar to those in QED [5]. A regular expression specifies a set of strings to be matched. It contains text characters (which match the corresponding characters in the strings being compared) and operator characters (which specify repetitions, choices, and other features). The letters of the alphabet and the digits are always text characters; thus the regular expression

```
integer
```

matches the string *integer* wherever it appears and the expression

```
a57D
```

looks for the string *a57D*.

Operators. The operator characters are

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

and if they are to be used as text characters, an escape should be used. The quotation mark operator (") indicates that whatever is contained between a pair of quotes is to be taken as text characters. Thus

```
xyz"++"
```

matches the string *xyz++* when it appears. Note that a part of a string may be quoted. It is harmless but unnecessary to quote an ordinary text character; the expression

```
"xyz++"
```

is the same as the one above. Thus by quoting every non-alphanumeric character being used as a text character, the user can avoid remembering the list above of current operator characters, and is safe should further extensions to Lex lengthen the list.

An operator character may also be turned into a text character by preceding it with \ as in

```
xyz\+\+
```

which is another, less readable, equivalent of the above

expressions. Another use of the quoting mechanism is to get a blank into an expression; normally, as explained above, blanks or tabs end a rule. Any blank character not contained within `[]` (see below) must be quoted. The usual C escapes with `\` are recognized: `\n` is newline, `\t` is tab, `\r` is return, and `\b` is backspace. To enter `\` itself, use `\\`. Since newline is illegal in an expression, `\n` must be used; it is not required to escape tab and backspace. Every character but blank, tab, newline and the list above is always a text character.

Character classes. Classes of characters can be specified using the operator pair `[]`. The construction `[ab]` matches a single character, which may be *a*, *b*, or *c*. Within square brackets, most operator meanings are ignored. Only three characters are special: these are `\`, `-` and `^`. The `-` character indicates ranges. For example,

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

indicates the character class containing all the lower case letters, the digits, the angle brackets, and underline. Ranges may be given in either order. Using `-` between any pair of characters which are not both upper case letters, both lower case letters, or both digits is implementation dependent and will get a warning message. (E.g., `[0-z]` in ASCII is many more characters than it is in EBCDIC). If it is desired to include the character `-` in a character class, it should be first or last; thus

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

matches all the digits and the two signs.

In character classes, the `^` operator must appear as the first character after the left bracket; it indicates that the resulting string is to be complemented with respect to the computer character set. Thus

```
[^abc]
```

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

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

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

Arbitrary character. To match almost any character, the operator character

is the class of all characters except newline. Escaping into octal is possible although non-portable:

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

matches all printable characters in the ASCII character set, from octal 40 (blank) to octal 176 (tilde).

Optional expressions. The operator `?` indicates an optional element of an expression. Thus

```
ab?c
```

matches either *ac* or *abc*.

Repeated expressions. Repetitions of classes are indicated by the operators `*` and `+`.

```
a*
```

is any number of consecutive *a* characters, including zero; while

```
a+
```

is one or more instances of *a*. For example,

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

is all strings of lower case letters. And

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

indicates all alphanumeric strings with a leading alphabetic character. This is a typical expression for recognizing identifiers in computer languages.

Alternation and Grouping. The operator `|` indicates alternation:

```
(ab|cd)
```

matches either *ab* or *cd*. Note that parentheses are used for grouping, although they are not necessary on the outside level;

```
ab|cd
```

would have sufficed. Parentheses can be used for more complex expressions:

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

matches such strings as *abefef*, *efefef*, *cdef*, or *cddd*; but not *abc*, *abcd*, or *abcdef*.

Context sensitivity. Lex will recognize a small amount of surrounding context. The two simplest operators for this are `^` and `$`. If the first character of an expression is `^`, the expression will only be matched at the beginning of a line (after a newline character, or at the beginning of the input stream). This can never conflict with the other meaning of `^`, complementation of character classes, since that only applies within the `[]` operators. If the very last character is `$`, the expression will only be matched at the end of a line (when immediately followed by newline). The latter operator is a special case of the `/` operator character, which indicates trailing context. The expression

```
ab/cd
```

matches the string *ab*, but only if followed by *cd*. Thus

```
ab$
```

is the same as

```
ab/\n
```

Left context is handled in Lex by *start conditions* as explained in section 10. If a rule is only to be executed when the Lex automaton interpreter is in start condition *x*, the rule should be prefixed by

```
<x>
```

using the angle bracket operator characters. If we considered "being at the beginning of a line" to be start condition *ONE*, then the \wedge operator would be equivalent to

```
<ONE>
```

Start conditions are explained more fully later.

Repetitions and Definitions. The operators $\{ \}$ specify either repetitions (if they enclose numbers) or definition expansion (if they enclose a name). For example

```
{digit}
```

looks for a predefined string named *digit* and inserts it at that point in the expression. The definitions are given in the first part of the Lex input, before the rules. In contrast,

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

looks for 1 to 5 occurrences of *a*.

Finally, initial $\%$ is special, being the separator for Lex source segments.

4 Lex Actions.

When an expression written as above is matched, Lex executes the corresponding action. This section describes some features of Lex which aid in writing actions. Note that there is a default action, which consists of copying the input to the output. This is performed on all strings not otherwise matched. Thus the Lex user who wishes to absorb the entire input, without producing any output, must provide rules to match everything. When Lex is being used with Yacc, this is the normal situation. One may consider that actions are what is done instead of copying the input to the output; thus, in general, a rule which merely copies can be omitted. Also, a character combination which is omitted from the rules and which appears as input is likely to be printed on the output, thus calling attention to the gap in the rules.

One of the simplest things that can be done is to ignore the input. Specifying a C null statement, $;$ as an action causes this result. A frequent rule is

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

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

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

```
" "
"\t"
"\n"
```

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

In more complex actions, the user will often want to know the actual text that matched some expression like $[a-z]^+$. Lex leaves this text in an external character array named *yytext*. Thus, to print the name found, a rule like

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

will print the string in *yytext*. The C function *printf* accepts a format argument and data to be printed; in this case, the format is "print string" ($\%$ indicating data conversion, and *s* indicating string type), and the data are the characters in *yytext*. So this just places the matched string on the output. This action is so common that it may be written as ECHO:

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

is the same as the above. Since the default action is just to print the characters found, one might ask why give a rule, like this one, which merely specifies the default action? Such rules are often required to avoid matching some other rule which is not desired. For example, if there is a rule which matches *read* it will normally match the instances of *read* contained in *bread* or *readjust*, to avoid this, a rule of the form $[a-z]^+$ is needed. This is explained further below.

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

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

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

```
yytext[yylen-1]
```

in C or

```
yytext(yylen)
```

in Ratfor.

Occasionally, a Lex action may decide that a rule has not recognized the correct span of characters. Two routines are provided to aid with this situation. First, *yymore()* can be called to indicate that the next input expression recognized is to be tacked on to the end of this input. Normally, the next input string would overwrite the current entry in *yytext*. Second, *yyless(n)* may be called to indicate that not all the characters matched by the currently successful expression are wanted right now. The argument *n* indicates the number of characters in *yytext* to be retained. Further characters previously matched are returned to the input. This provides the same sort of lookahead offered by the / operator, but in a different form.

Example: Consider a language which defines a string as a set of characters between quotation (") marks, and provides that to include a " in a string it must be preceded by a \. The regular expression which matches that is somewhat confusing, so that it might be preferable to write

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

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

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

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

which prints a message, returns the letter after the operator to the input stream, and treats the operator as "--". Alternatively it might be desired to treat this as "= -a". To do this, just return the minus sign as well as the letter to the input:

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

will perform the other interpretation. Note that the expressions for the two cases might more easily be written

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

in the first case and

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

in the second; no backup would be required in the rule action. It is not necessary to recognize the whole identifier to observe the ambiguity. The possibility of "--3", however, makes

```
==/[^\t\n]
```

a still better rule.

In addition to these routines, Lex also permits access to the I/O routines it uses. They are:

- 1) *input()* which returns the next input character;
- 2) *output(c)* which writes the character *c* on the output; and
- 3) *unput(c)* pushes the character *c* back onto the input stream to be read later by *input()*.

By default these routines are provided as macro definitions, but the user can override them and supply private versions. There is another important routine in Ratfor, named *lexshf*, which is described below under "Character Set". These routines define the relationship between external files and internal characters, and must all be retained or modified consistently. They may be redefined, to cause input or output to be transmitted to or from strange places, including other programs or internal memory; but the character set used must be consistent in all routines; a value of zero returned by *input* must mean end of file; and the relationship between *unput* and *input* must be retained or the Lex lookahead will not work. Lex does not look ahead at all if it does not have to, but every rule ending in + * ? or \$ or containing / implies lookahead. Lookahead is also necessary to match an expression that is a prefix of another expression. See below for a discussion of the character set used by Lex. The standard Lex library imposes a 100 character limit on backup.

Another Lex library routine that the user will sometimes want to redefine is *yywrap()* which is called whenever Lex reaches an end-of-file. If *yywrap* returns a 1, Lex continues with the normal wrapup on end of input. Sometimes, however, it is convenient to arrange for more input to arrive from a new source. In this case, the user should provide a *yywrap* which arranges for new input and returns 0. This instructs Lex to continue processing. The default *yywrap* always returns 1.

This routine is also a convenient place to print tables, summaries, etc. at the end of a program. Note that it is not possible to write a normal rule which recognizes end-of-file; the only access to this condition is through *yywrap*. In fact, unless a private version of *input()* is supplied a file containing nulls cannot be handled, since a value of 0 returned by *input* is taken to be end-of-file.

In Ratfor all of the standard I/O library routines, *input*,

output, *unput*, *yywrap*, and *lexshf*, are defined as integer functions. This requires *input* and *yywrap* to be called with arguments. One dummy argument is supplied and ignored.

5 Ambiguous Source Rules.

Lex can handle ambiguous specifications. When more than one expression can match the current input, Lex chooses as follows:

- 1) The longest match is preferred.
- 2) Among rules which matched the same number of characters, the rule given first is preferred.

Thus, suppose the rules

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

to be given in that order. If the input is *integers*, it is taken as an identifier, because *[a-z]+* matches 8 characters while *integer* matches only 7. If the input is *integer*, both rules match 7 characters, and the keyword rule is selected because it was given first. Anything shorter (e.g. *int*) will not match the expression *integer* and so the identifier interpretation is used.

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

*.**

might seem a good way of recognizing a string in single quotes. But it is an invitation for the program to read far ahead, looking for a distant single quote. Presented with the input

'first' quoted string here, 'second' here

the above expression will match

'first' quoted string here, 'second'

which is probably not what was wanted. A better rule is of the form

'[^\n]*'

which, on the above input, will stop after *'first'*. The consequences of errors like this are mitigated by the fact that the *.* operator will not match newline. Thus expressions like *.* stop* on the current line. Don't try to defeat this with expressions like *[^\n]+* or equivalents; the Lex generated program will try to read the entire input file, causing internal buffer overflows.

Note that Lex is normally partitioning the input stream, not searching for all possible matches of each expression. This means that each character is accounted for once and only once. For example, suppose it is desired to count occurrences of both *she* and *he* in an input text. Some Lex rules to do this might be

```
she s++ ;
he h++ ;
\n |
;
```

where the last two rules ignore everything besides *he* and *she*. Remember that *.* does not include newline. Since *she* includes *he*, Lex will normally *not* recognize the instances of *he* included in *she*, since once it has passed a *she* those characters are gone.

Sometimes the user would like to override this choice. The action REJECT means "go do the next alternative." It causes whatever rule was second choice after the current rule to be executed. The position of the input pointer is adjusted accordingly. Suppose the user really wants to count the included instances of *he*:

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

these rules are one way of changing the previous example to do just that. After counting each expression, it is rejected; whenever appropriate, the other expression will then be counted. In this example, of course, the user could note that *she* includes *he* but not vice versa, and omit the REJECT action on *he*; in other cases, however, it would not be possible a priori to tell which input characters were in both classes.

Consider the two rules

```
a[bc]+ { ... ; REJECT;}
a[cd]+ { ... ; REJECT;};
```

If the input is *ab*, only the first rule matches, and on *ad* only the second matches. The input string *accb* matches the first rule for four characters and then the second rule for three characters. In contrast, the input *accd* agrees with the second rule for four characters and then the first rule for three.

In general, REJECT is useful whenever the purpose of Lex is not to partition the input stream but to detect all examples of some items in the input, and the instances of these items may overlap or include each other. Suppose a digram table of the input is desired; normally the digrams overlap, that is the word *the* is considered to contain both *th* and *he*. Assuming a two-dimensional array named *digram* to be incremented, the appropriate source is

```
%%
[a-z][a-z] {digram[yytext[0]][yytext[1]]++ ; REJECT;}
. ;
\n ;
```

where the REJECT is necessary to pick up a letter pair beginning at every character, rather than at every other character.

6 Lex Source Definitions.

Remember the format of the Lex source:

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

So far only the rules have been described. The user needs additional options, though, to define variables for use in his program and for use by Lex. These can go either in the definitions section or in the rules section.

Remember that Lex is turning the rules into a program. Any source not intercepted by Lex is copied into the generated program. There are three classes of such things.

- 1) Any line which is not part of a Lex rule or action which begins with a blank or tab is copied into the Lex generated program. Such source input prior to the first %% delimiter will be external to any function in the code; if it appears immediately after the first %%, it appears in an appropriate place for declarations in the function written by Lex which contains the actions. This material must look like program fragments, and should precede the first Lex rule.

As a side effect of the above, lines which begin with a blank or tab, and which contain a comment, are passed through to the generated program. This can be used to include comments in either the Lex source or the generated code. The comments should follow the host language convention.

- 2) Anything included between lines containing only %{ and %} is copied out as above. The delimiters are discarded. This format permits entering text like preprocessor statements that must begin in column 1, or copying lines that do not look like programs.
- 3) Anything after the third %% delimiter, regardless of formats, etc., is copied out after the Lex output.

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

```
name translation
```

and it causes the string given as a translation to be associated with the name. The name and translation must be separated by at least one blank or tab, and the name must begin with a letter. The translation can then be called out by the {name} syntax in a rule. Using {D} for the digits and {E} for an exponent field, for example, might abbreviate rules to recognize numbers:

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

Note the first two rules for real numbers; both require a decimal point and contain an optional exponent field, but the first requires at least one digit before the decimal point and the second requires at least one digit after the decimal point. To correctly handle the problem posed by a Fortran expression such as *35.EQ.1*, which does not contain a real number, a context-sensitive rule such as

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

could be used in addition to the normal rule for integers.

The definitions section may also contain other commands, including the selection of a host language, a character set table, a list of start conditions, or adjustments to the default size of arrays within Lex itself for larger source programs. These possibilities are discussed below under "Summary of Source Format," section 12.

7 Usage.

There are two steps in compiling a Lex source program. First, the Lex source must be turned into a generated program in the host general purpose language. Then this program must be compiled and loaded, usually with a library of Lex subroutines. The generated program is on a file named *lex.yy.c* for a C host language source and *lex.yy.r* for a Ratfor host environment. There are two I/O libraries, one for C defined in terms of the C standard library [6], and the other defined in terms of Ratfor. To indicate that a Lex source file is intended to be used with the Ratfor host language, make the first line of the file %R.

The C programs generated by Lex are slightly different on OS/370, because the OS compiler is less powerful than the UNIX or GCOS compilers, and does less at compile time. C programs generated on GCOS and UNIX are the same. The C host language is default, but may be explicitly requested by making the first line of the source file %C.

The Ratfor generated by Lex is the same on all systems, but can not be compiled directly on TSO. See below for instructions. The Ratfor I/O library, however, varies slightly because the different Fortrans disagree on the method of indicating end-of-input and the name of the library routine for logical AND. The Ratfor I/O library, dependent on Fortran character I/O, is quite slow. In particular it reads all input lines as 80A1 format; this will truncate any longer line, discarding your data, and pads any shorter line with blanks. The library version of *input* removes the padding (including any trailing blanks from the original input) before processing. Each source

file using a Ratfor host should begin with the "%R" command.

UNIX. The libraries are accessed by the loader flags *-llc* for C and *-llr* for Ratfor; the C name may be abbreviated to *-ll*. So an appropriate set of commands is

C Host	Ratfor Host
lex source	lex source
cc lex.yy.c -ll -lS	rc -2 lex.yy.r -llr

The resulting program is placed on the usual file *a.out* for later execution. To use Lex with Yacc see below. Although the default Lex I/O routines use the C standard library, the Lex automata themselves do not do so; if private versions of *input*, *output* and *unput* are given, the library can be avoided. Note the "-2" option in the Ratfor compile command; this requests the larger version of the compiler, a useful precaution.

GCOS. The Lex commands on GCOS are stored in the ".l" library. The appropriate command sequences are:

C Host	Ratfor Host
./lex source	./lex source
./cc lex.yy.c ./lexlib h =	./rc a = lex.yy.r ./lexlib h =

The resulting program is placed on the usual file *.program* for later execution (as indicated by the "h=" option); it may be copied to a permanent file if desired. Note the "a=" option in the Ratfor compile command; this indicates that the Fortran compiler is to run in ASCII mode.

TSO. Lex is just barely available on TSO. Restrictions imposed by the compilers which must be used with its output make it rather inconvenient. To use the C version, type

```
exec 'dot.lex.clist(lex)' 'sourcename'
exec 'dot.lex.clist(rlload)' 'libraryname membername'
```

The first command analyzes the source file and writes a C program on file *lex.yy.text*. The second command runs this file through the C compiler and links it with the Lex C library (stored on 'hr289.lcl.load') placing the object program in your file *libraryname.LOAD(membername)* as a completely linked load module. The compiling command uses a special version of the C compiler command on TSO which provides an unusually large intermediate assembler file to compensate for the unusual bulk of C-compiled Lex programs on the OS system. Even so, almost any Lex source program is too big to compile, and must be split.

The same Lex command will compile Ratfor Lex programs, leaving a file *lex.yy.rat* instead of *lex.yy.text* in your directory. The Ratfor program must be edited, however, to compensate for peculiarities of IBM Ratfor. A command sequence to do this, and then compile and load, is available. The full commands are:

```
exec 'dot.lex.clist(lex)' 'sourcename'
```

```
exec 'dot.lex.clist(rlload)' 'libraryname membername'
```

with the same overall effect as the C language commands. However, the Ratfor commands will run in a 150K byte partition, while the C commands require 250K bytes to operate.

The steps involved in processing the generated Ratfor program are:

- a. Edit the Ratfor program.
 1. Remove all tabs.
 2. Change all lower case letters to upper case letters.
 3. Convert the file to an 80-column card image file.
- b. Process the Ratfor through the Ratfor preprocessor to get Fortran code.
- c. Compile the Fortran.
- d. Load with the libraries 'hr289.lrl.load' and 'sys1.fortlib'.

The final load module will only read input in 80-character fixed length records. **Warning:** Work is in progress on the IBM C compiler, and Lex and its availability on the IBM 370 are subject to change without notice.

8 Lex and Yacc.

If you want to use Lex with Yacc, note that what Lex writes is a program named *yylex()*, the name required by Yacc for its analyzer. Normally, the default main program on the Lex library calls this routine, but if Yacc is loaded, and its main program is used, Yacc will call *yylex()*. In this case each Lex rule should end with

```
return(token);
```

where the appropriate token value is returned. An easy way to get access to Yacc's names for tokens is to compile the Lex output file as part of the Yacc output file by placing the line

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

in the last section of Yacc input. Supposing the grammar to be named "good" and the lexical rules to be named "better" the UNIX command sequence can just be:

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

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

9 Examples.

As a trivial problem, consider copying an input file while adding 3 to every positive number divisible by 7. Here is a suitable Lex source program

```

%%
int k;
[0-9]+ {
scanf(-1, yytext, "%d", &k);
if (k%7 == 0)
printf("%d", k+3);
else
printf("%d", k);
}

```

to do just that. The rule `[0-9]+` recognizes strings of digits; `scanf` converts the digits to binary and stores the result in `k`. The operator `%` (remainder) is used to check whether `k` is divisible by 7; if it is, it is incremented by 3 as it is written out. It may be objected that this program will alter such input items as `49.63` or `X7`. Furthermore, it increments the absolute value of all negative numbers divisible by 7. To avoid this, just add a few more rules after the active one, as here:

```

%%
-?[0-9]+ int k;
{
scanf(-1, yytext, "%d", &k);
printf("%d", k%7 == 0 ? k+3 : k);
}
-?[0-9.]+ ECHO;
[A-Za-z][A-Za-z0-9]+ ECHO;

```

Numerical strings containing a `.` or preceded by a letter will be picked up by one of the last two rules, and not changed. The *if-else* has been replaced by a C conditional expression to save space; the form `a?b:c` means "if `a` then `b` else `c`".

For an example of statistics gathering, here is a program which histograms the lengths of words, where a word is defined as a string of letters.

```

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

```

This program accumulates the histogram, while producing no output. At the end of the input it prints the table. The final statement `return(1)`; indicates that Lex is to perform wrapup. If `yywrap` returns zero (false) it implies that further input is available and the program is to continue reading and processing. To provide a `yywrap` that

never returns true causes an infinite loop.

As a larger example, here are some parts of a program written by N. L. Schryer to convert double precision Fortran to single precision Fortran. Because Fortran does not distinguish upper and lower case letters, this routine begins by defining a set of classes including both cases of each letter:

```

a [aA]
b [bB]
c [cC]
...
z [zZ]

```

An additional class recognizes white space:

```
W [\t]*
```

The first rule changes "double precision" to "real", or "DOUBLE PRECISION" to "REAL".

```
{d}{o}{u}{b}{l}{e}{W}{p}{r}{e}{c}{i}{s}{i}{o}{n} {
printf(yytext[0] == 'd'? "real" : "REAL");
}

```

Care is taken throughout this program to preserve the case (upper or lower) of the original program. The conditional operator is used to select the proper form of the keyword. The next rule copies continuation card indications to avoid confusing them with constants:

```
" "[^ 0] ECHO;
```

In the regular expression, the quotes surround the blanks. It is interpreted as "beginning of line, then five blanks, then anything but blank or zero." Note the two different meanings of `^`. There follow some rules to change double precision constants to ordinary floating constants.

```
[0-9]+{W}{d}{W}{+}?{W}[0-9]+ |
[0-9]+{W}."{W}{d}{W}{+}?{W}[0-9]+ |
"."{W}[0-9]+{W}{d}{W}{+}?{W}[0-9]+ {
/* convert constants */
for(p=yytext; *p != 0; p++)
{
if (*p == 'd' | *p == 'D')
*p = + 'e' - 'd';
ECHO;
}
}

```

After the floating point constant is recognized, it is scanned by the `for` loop to find the letter `d` or `D`. The program then adds `'e'-'d'`, which converts it to the next letter of the alphabet. The modified constant, now single-precision, is written out again. There follow a series of names which must be respelled to remove their initial `d`. By using the array `yytext` the same action suffices for all the names (only a sample of a rather long list is given here).


```
{d}{s}{r}{n} |
{d}{c}{o}{s} |
{d}{s}{q}{r}{t} |
{d}{a}{t}{a}{n} |
...
{d}{f}{t}{o}{a}{t} printf("%s",yytext+1);
```

Another list of names must have initial *d* changed to initial *a*:

```
{d}{l}{o}{g} |
{d}{l}{o}{g}10 |
{d}{m}{i}{n}1 |
{d}{m}{a}{x}1 |
yytext[0] = + 'a' - 'd';
ECHO;
}
```

And one routine must have initial *d* changed to initial *r*:

```
{d}l{m}{a}{c}{h} {yytext[0] = + 'r' - 'd';
```

To avoid such names as *dsinx* being detected as instances of *dsin*, some final rules pick up longer words as identifiers and copy some surviving characters:

```
[A-Za-z][A-Za-z0-9]* |
[0-9]+ |
\n |
ECHO;
```

Note that this program is not complete; it does not deal with the spacing problems in Fortran or with the use of keywords as identifiers.

10 Left Context Sensitivity.

Sometimes it is desirable to have several sets of lexical rules to be applied at different times in the input. For example, a compiler preprocessor might distinguish preprocessor statements and analyze them differently from ordinary statements. This requires sensitivity to prior context, and there are several ways of handling such problems. The *^* operator, for example, is a prior context operator, recognizing immediately preceding left context just as *\$* recognizes immediately following right context. Adjacent left context could be extended, to produce a facility similar to that for adjacent right context, but it is unlikely to be as useful, since often the relevant left context appeared some time earlier, such as at the beginning of a line.

This section describes three means of dealing with different environments: a simple use of flags, when only a few rules change from one environment to another, the use of *start conditions* on rules, and the possibility of making multiple lexical analyzers all run together. In each case, there are rules which recognize the need to change the environment in which the following input text

is analyzed, and set some parameter to reflect the change. This may be a flag explicitly tested by the user's action code; such a flag is the simplest way of dealing with the problem, since Lex is not involved at all. It may be **more** convenient, however, to have Lex remember the flags as initial conditions on the rules. Any rule may be associated with a start condition. It will only be recognized when Lex is in that start condition. The current start condition may be changed at any time. Finally, if the sets of rules for the different environments are very dissimilar, clarity may be best achieved by writing several distinct lexical analyzers, and switching from one to another as desired.

Consider the following problem: copy the input to the output, changing the word *magic* to *first* on every line which began with the letter *a*, changing *magic* to *second* on every line which began with the letter *b*, and changing *magic* to *third* on every line which began with the letter *c*. All other words and all other lines are left unchanged.

These rules are so simple that the easiest way to do this job is with a flag:

```
int flag;

%%
^a {flag = 'a'; ECHO;}
^b {flag = 'b'; ECHO;}
^c {flag = 'c'; ECHO;}
\n {flag = 0 ; ECHO;}
magic {
switch (flag)
{
case 'a': printf("first"); break;
case 'b': printf("second"); break;
case 'c': printf("third"); break;
default: ECHO; break;
}
}
```

should be adequate.

To handle the same problem with start conditions, each start condition must be introduced to Lex in the definitions section with a line reading

```
%Start name1 name2 ...
```

where the conditions may be named in any order. The word *Start* may be abbreviated to *s* or *S*. The conditions may be referenced at the head of a rule with the *<>* brackets:

```
<name1>expression
```

is a rule which is only recognized when Lex is in the start condition *name1*. To enter a start condition, execute the action statement

```
BEGIN name1;
```

which changes the start condition to *name1*. To resume the normal state,

BEGIN 0;

resets the initial condition of the Lex automaton interpreter. A rule may be active in several start conditions:

<name1,name2,name3>

is a legal prefix. Any rule not beginning with the <> prefix operator is always active.

The same example as before can be written:

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

where the logic is exactly the same as in the previous method of handling the problem, but Lex does the work rather than the user's code.

11 Character Set.

The programs generated by Lex handle character I/O only through the routines *input*, *output*, and *unput*. Thus the character representation provided in these routines is accepted by Lex and employed to return values in *yytext*. For internal use a character is represented as a small integer which, if the standard library is used, has a value equal to the integer value of the bit pattern representing the character on the host computer. In C, the I/O routines are assumed to deal directly in this representation. In Ratfor, it is anticipated that many users will prefer left-adjusted rather than right-adjusted characters; thus the routine *lexshf* is called to change the representation delivered by *input* into a right-adjusted integer. If the user changes the I/O library, the routine *lexshf* should also be changed to a compatible version. The Ratfor library I/O system is arranged to represent the letter *a* as in the Fortran value *IHa* while in C the letter *a* is represented as the character constant '*a*'. If this interpretation is changed, by providing I/O routines which translate the characters, Lex must be told about it, by giving a translation table. This table must be in the definitions section, and must be bracketed by lines containing only "%T". The table contains lines of the form

{integer} {character string}

which indicate the value associated with each character. Thus the next example maps the lower and upper case letters together into the integers 1 through 26, newline into 27, + and - into 28 and 29, and the digits into 30 through 39. Note the escape for newline. If a table is supplied, every character that is to appear either in the

```
%T
1  Aa
2  Bb
...
26 Zz
27 \n
28 +
29 -
30 0
31 1
...
39 9
%T
```

Sample character table.

rules or in any valid input must be included in the table. No character may be assigned the number 0, and no character may be assigned a bigger number than the size of the hardware character set.

It is not likely that C users will wish to use the character table feature; but for Fortran portability it may be essential.

Although the contents of the Lex Ratfor library routines for input and output run almost unmodified on UNIX, GCOS, and OS/370, they are not really machine independent, and would not work with CDC or Burroughs Fortran compilers. The user is of course welcome to replace *input*, *output*, *unput* and *lexshf* but to replace them by completely portable Fortran routines is likely to cause a substantial decrease in the speed of Lex Ratfor programs. A simple way to produce portable routines would be to leave *input* and *output* as routines that read with 80A1 format, but replace *lexshf* by a table lookup routine.

12 Summary of Source Format.

The general form of a Lex source file is:

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

The definitions section contains a combination of

- 1) Definitions, in the form "name space translation".
- 2) Included code, in the form "space code".
- 3) Included code, in the form

```
%{
code
%}
```

- 4) Start conditions, given in the form

```
%S name1 name2 ...
```

- 5) Character set tables, in the form

```
%T
number space character-string
...
%T
```

- 6) A language specifier, which must also precede any rules or included code, in the form "%C" for C or "%R" for Ratfor.
- 7) Changes to internal array sizes, in the form

```
%x nnn
```

where *nnn* is a decimal integer representing an array size and *x* selects the parameter as follows:

Letter	Parameter
p	positions
n	states
e	tree nodes
a	transitions
k	packed character classes
o	output array size

Lines in the rules section have the form "expression action" where the action may be continued on succeeding lines by using braces to delimit it.

Regular expressions in Lex use the following operators:

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

13 Caveats and Bugs.

There are pathological expressions which produce exponential growth of the tables when converted to deterministic machines; fortunately, they are rare.

REJECT does not rescan the input; instead it remembers the results of the previous scan. This means that if a rule with trailing context is found, and REJECT executed, the user must not have used *unput* to change the characters forthcoming from the input stream. This is the only restriction on the user's ability to manipulate the not-yet-processed input.

TSO Lex is an older version. Among the non-supported features are REJECT, start conditions, or variable length trailing context, And any significant Lex source is too big for the IBM C compiler when translated.

14 Acknowledgments.


As should be obvious from the above, the outside of Lex is patterned on Yacc and the inside on Aho's string matching routines. Therefore, both S. C. Johnson and A. V. Aho are really originators of much of Lex, as well as debuggers of it. Many thanks are due to both.


The code of the current version of Lex was designed, written, and debugged by Eric Schmidt.

15 References.

1. D. M. Ritchie, B. W. Kernighan, and M. E. Lesk, *The C Programming Language*, Computing Science Technical Report No. 31 (1975). Bell Laboratories, Murray Hill, NJ 07974.
2. B. W. Kernighan, *Ratfor: A Preprocessor for a Rational Fortran*, to appear in *Software Practice and Experience*, 1975.
3. S. C. Johnson, *Yacc: Yet Another Compiler Compiler*, Computing Science Technical Report No. 32, 1975, Bell Laboratories, Murray Hill, NJ 07974.
4. A. V. Aho and M. J. Corasick, *Efficient String Matching: An Aid to Bibliographic Search*, Comm. ACM 18, 333-340 (1975).
5. B. W. Kernighan, D. M. Ritchie and K. L. Thompson, *QED Text Editor*, Computing Science Technical Report No. 5, 1972, Bell Laboratories, Murray Hill, NJ 07974.
6. D. M. Ritchie, private communication. See also M. E. Lesk, *The Portable C Library*, contained in reference [1], above.







RATFOR — A Preprocessor for a Rational Fortran

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

Although Fortran is not a pleasant language to use, it does have the advantages of universality and (usually) relative efficiency. The Ratfor language attempts to conceal the main deficiencies of Fortran while retaining its desirable qualities, by providing decent control flow statements:

- statement grouping
- **if-else** and **switch** for decision-making
- **while**, **for**, **do**, and **repeat-until** for looping
- **break** and **next** for controlling loop exits

and some "syntactic sugar":

- free form input (multiple statements/line, automatic continuation)
- unobtrusive comment convention
- translation of $>$, $> =$, etc., into `.GT.`, `.GE.`, etc.
- **return**(expression) statement for functions
- **define** statement for symbolic parameters
- **include** statement for including source files

Ratfor is implemented as a preprocessor which translates this language into Fortran.

Once the control flow and cosmetic deficiencies of Fortran are hidden, the resulting language is remarkably pleasant to use. Ratfor programs are markedly easier to write, and to read, and thus easier to debug, maintain and modify than their Fortran equivalents.

It is readily possible to write Ratfor programs which are portable to other environments. Ratfor is written in itself in this way, so it is also portable; versions of Ratfor are now running on at least a dozen different types of computers at over one hundred locations.

This paper discusses design criteria for a Fortran preprocessor, the Ratfor language and its implementation, and user experience.

January 1, 1977



RATFOR — A Preprocessor for a Rational Fortran

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

1. INTRODUCTION

Most programmers will agree that Fortran is an unpleasant language to program in, yet there are many occasions when they are forced to use it. For example, Fortran is often the only language thoroughly supported on the local computer. Indeed, it is the closest thing to a universal programming language currently available: with care it is possible to write large, truly portable Fortran programs[1]. Finally, Fortran is often the most "efficient" language available, particularly for programs requiring much computation.

But Fortran *is* unpleasant. Perhaps the worst deficiency is in the control flow statements — conditional branches and loops — which express the logic of the program. The conditional statements in Fortran are primitive. The Arithmetic IF forces the user into at least two statement numbers and two (implied) GOTO's; it leads to unintelligible code, and is eschewed by good programmers. The Logical IF is better, in that the test part can be stated clearly, but hopelessly restrictive because the statement that follows the IF can only be one Fortran statement (with some *further* restrictions!). And of course there can be no ELSE part to a Fortran IF: there is no way to specify an alternative action if the IF is not satisfied.

The Fortran DO restricts the user to going forward in an arithmetic progression. It is fine for "1 to N in steps of 1 (or 2 or ...)"; but there is no direct way to go backwards, or even (in ANSI Fortran[2]) to go from 1 to N-1. And of course the DO is useless if one's problem doesn't map into an arithmetic progression.

The result of these failings is that Fortran programs must be written with numerous labels and branches. The resulting code is particularly difficult to read and understand, and thus hard to debug and modify.

When one is faced with an unpleasant language, a useful technique is to define a new language that overcomes the deficiencies, and to translate it into the unpleasant one with a preprocessor. This is the approach taken with

Ratfor. (The preprocessor idea is of course not new, and preprocessors for Fortran are especially popular today. A recent listing [3] of preprocessors shows more than 50, of which at least half a dozen are widely available.)

2. LANGUAGE DESCRIPTION

Design

Ratfor attempts to retain the merits of Fortran (universality, portability, efficiency) while hiding the worst Fortran inadequacies. The language *is* Fortran except for two aspects. First, since control flow is central to any program, regardless of the specific application, the primary task of Ratfor is to conceal this part of Fortran from the user, by providing decent control flow structures. These structures are sufficient and comfortable for structured programming in the narrow sense of programming without GOTO's. Second, since the preprocessor must examine an entire program to translate the control structure, it is possible at the same time to clean up many of the "cosmetic" deficiencies of Fortran, and thus provide a language which is easier and more pleasant to read and write.

Beyond these two aspects — control flow and cosmetics — Ratfor does nothing about the host of other weaknesses of Fortran. Although it would be straightforward to extend it to provide character strings, for example, they are not needed by everyone, and of course the preprocessor would be harder to implement. Throughout, the design principle which has determined what should be in Ratfor and what should not has been *Ratfor doesn't know any Fortran*. Any language feature which would require that Ratfor really understand Fortran has been omitted. We will return to this point in the section on implementation.

Even within the confines of control flow and cosmetics, we have attempted to be selective in what features to provide. The intent has been to provide a small set of the most useful constructs, rather than to throw in everything that has ever been thought useful by someone.

The rest of this section contains an informal description of the Ratfor language. The control flow aspects will be quite familiar to readers used to languages like Algol, PL/I, Pascal, etc., and the cosmetic changes are equally straightforward. We shall concentrate on showing what the language looks like.

Statement Grouping

Fortran provides no way to group statements together, short of making them into a subroutine. The standard construction "if a condition is true, do this group of things," for example,

```
if (x > 100)
    { call error("x>100"); err = 1; return }
```

cannot be written directly in Fortran. Instead a programmer is forced to translate this relatively clear thought into murky Fortran, by stating the negative condition and branching around the group of statements:

```
if (x .le. 100) goto 10
    call error(5hx>100)
    err = 1
    return
10    ...
```

When the program doesn't work, or when it must be modified, this must be translated back into a clearer form before one can be sure what it does.

Ratfor eliminates this error-prone and confusing back-and-forth translation: the first form is the way the computation is written in Ratfor. A group of statements can be treated as a unit by enclosing them in the braces { and }. This is true throughout the language: wherever a single Ratfor statement can be used, there can be several enclosed in braces. (Braces seem clearer and less obtrusive than **begin** and **end** or **do** and **end**, and of course **do** and **end** already have Fortran meanings.)

Cosmetics contribute to the readability of code, and thus to its understandability. The character ">" is clearer than ".GT.", so Ratfor translates it appropriately, along with several other similar shorthands. Although many Fortran compilers permit character strings in quotes (like "x>100"), quotes are not allowed in ANSI Fortran, so Ratfor converts it into the right number of H's: computers count better than people do.

Ratfor is a free-form language: statements may appear anywhere on a line, and several may appear on one line if they are separated by semi-

colons. The example above could also be written as

```
if (x > 100) {
    call error("x>100")
    err = 1
    return
}
```

In this case, no semicolon is needed at the end of each line because Ratfor assumes there is one statement per line unless told otherwise.

Of course, if the statement that follows the **if** is a single statement (Ratfor or otherwise), no braces are needed:

```
if (y <= 0.0 & z <= 0.0)
    write(6, 20) y, z
```

No continuation need be indicated because the statement is clearly not finished on the first line. In general Ratfor continues lines when it seems obvious that they are not yet done. (The continuation convention is discussed in detail later.)

Although a free-form language permits wide latitude in formatting styles, it is wise to pick one that is readable, then stick to it. In particular, proper indentation is vital, to make the logical structure of the program obvious to the reader.

The "else" Clause

Ratfor provides an **else** statement to handle the construction "if a condition is true, do this thing, otherwise do that thing."

```
if (a <= b)
    { sw = 0; write(6, 1) a, b }
else
    { sw = 1; write(6, 1) b, a }
```

This writes out the smaller of **a** and **b**, then the larger, and sets **sw** appropriately.

The Fortran equivalent of this code is circuitous indeed:

```
if (a .gt. b) goto 10
    sw = 0
    write(6, 1) a, b
    goto 20
10    sw = 1
    write(6, 1) b, a
20    ...
```

This is a mechanical translation; shorter forms exist, as they do for many similar situations. But all translations suffer from the same problem: since they are translations, they are less clear and understandable than code that is not a transla-

tion. To understand the Fortran version, one must scan the entire program to make sure that no other statement branches to statements 10 or 20 before one knows that indeed this is an **if-else** construction. With the Ratfor version, there is no question about how one gets to the parts of the statement. The **if-else** is a single unit, which can be read, understood, and ignored if not relevant. The program says what it means.

As before, if the statement following an **if** or an **else** is a single statement, no braces are needed:

```

if (a <= b)
    sw = 0
else
    sw = 1

```

The syntax of the **if** statement is

```

if (legal Fortran condition)
    Ratfor statement
else
    Ratfor statement

```

where the **else** part is optional. The *legal Fortran condition* is anything that can legally go into a Fortran Logical IF. Ratfor does not check this clause, since it does not know enough Fortran to know what is permitted. The *Ratfor statement* is any Ratfor or Fortran statement, or any collection of them in braces.

Nested if's

Since the statement that follows an **if** or an **else** can be any Ratfor statement, this leads immediately to the possibility of another **if** or **else**. As a useful example, consider this problem: the variable **f** is to be set to **-1** if **x** is less than zero, to **+1** if **x** is greater than 100, and to **0** otherwise. Then in Ratfor, we write

```

if (x < 0)
    f = -1
else if (x > 100)
    f = +1
else
    f = 0

```

Here the statement after the first **else** is another **if-else**. Logically it is just a single statement, although it is rather complicated.

This code says what it means. Any version written in straight Fortran will necessarily be indirect because Fortran does not let you say what you mean. And as always, clever shortcuts may turn out to be too clever to understand a year from now.

Following an **else** with an **if** is one way to write a multi-way branch in Ratfor. In general the structure

```

if (...)
    ---
else if (...)
    ---
else if (...)
    ---
...
else
    ---

```

provides a way to specify the choice of exactly one of several alternatives. (Ratfor also provides a **switch** statement which does the same job in certain special cases; in more general situations, we have to make do with spare parts.) The tests are laid out in sequence, and each one is followed by the code associated with it. Read down the list of decisions until one is found that is satisfied. The code associated with this condition is executed, and then the entire structure is finished. The trailing **else** part handles the "default" case, where none of the other conditions apply. If there is no default action, this final **else** part is omitted:

```

if (x < 0)
    x = 0
else if (x > 100)
    x = 100

```

if-else ambiguity

There is one thing to notice about complicated structures involving nested **if's** and **else's**. Consider

```

if (x > 0)
    if (y > 0)
        write(6, 1) x, y
    else
        write(6, 2) y

```

There are two **if's** and only one **else**. Which **if** does the **else** go with?

This is a genuine ambiguity in Ratfor, as it is in many other programming languages. The ambiguity is resolved in Ratfor (as elsewhere) by saying that in such cases the **else** goes with the closest previous un-**else'd if**. Thus in this case, the **else** goes with the inner **if**, as we have indicated by the indentation.

It is a wise practice to resolve such cases by explicit braces, just to make your intent clear. In the case above, we would write

```

if (x > 0) {
    if (y > 0)
        write(6, 1) x, y
    else
        write(6, 2) y
}

```

which does not change the meaning, but leaves no doubt in the reader's mind. If we want the other association, we *must* write

```

if (x > 0) {
    if (y > 0)
        write(6, 1) x, y
}
else
    write(6, 2) y

```

The "switch" Statement

The **switch** statement provides a clean way to express multi-way branches which branch on the value of some integer-valued expression. The syntax is

```

switch (expression) {
    case expr1 :
        statements
    case expr2, expr3 :
        statements
    ...
    default:
        statements
}

```

Each **case** is followed by a list of comma-separated integer expressions. The *expression* inside **switch** is compared against the case expressions *expr1*, *expr2*, and so on in turn until one matches, at which time the statements following that **case** are executed. If no cases match *expression*, and there is a **default** section, the statements with it are done; if there is no **default**, nothing is done. In all situations, as soon as some block of statements is executed, the entire **switch** is exited immediately. (Readers familiar with C[4] should beware that this behavior is not the same as the C **switch**.)

The "do" Statement

The **do** statement in Ratfor is quite similar to the **DO** statement in Fortran, except that it uses no statement number. The statement number, after all, serves only to mark the end of the **DO**, and this can be done just as easily with braces. Thus

```

do i = 1, n {
    x(i) = 0.0
    y(i) = 0.0
    z(i) = 0.0
}

```

is the same as

```

do 10 i = 1, n
    x(i) = 0.0
    y(i) = 0.0
    z(i) = 0.0
10 continue

```

The syntax is:

```

do legal-Fortran-DO-text
    Ratfor statement

```

The part that follows the keyword **do** has to be something that can legally go into a Fortran **DO** statement. Thus if a local version of Fortran allows **DO** limits to be expressions (which is not currently permitted in ANSI Fortran), they can be used in a Ratfor **do**.

The *Ratfor statement* part will often be enclosed in braces, but as with the **if**, a single statement need not have braces around it. This code sets an array to zero:

```

do i = 1, n
    x(i) = 0.0

```

Slightly more complicated,

```

do i = 1, n
    do j = 1, n
        m(i, j) = 0

```

sets the entire array **m** to zero, and

```

do i = 1, n
    do j = 1, n
        if (i < j)
            m(i, j) = -1
        else if (i == j)
            m(i, j) = 0
        else
            m(i, j) = +1

```

sets the upper triangle of **m** to -1 , the diagonal to zero, and the lower triangle to $+1$. (The operator **==** is "equals", that is, "EQ.") In each case, the statement that follows the **do** is logically a *single* statement, even though complicated, and thus needs no braces.

"break" and "next"

Ratfor provides a statement for leaving a loop early, and one for beginning the next iteration. **break** causes an immediate exit from the

do; in effect it is a branch to the statement *after* the **do**. **next** is a branch to the bottom of the loop, so it causes the next iteration to be done. For example, this code skips over negative values in an array:

```

do i = 1, n {
  if (x(i) < 0.0)
    next
  process positive element
}

```

break and **next** also work in the other Ratfor looping constructions that we will talk about in the next few sections.

break and **next** can be followed by an integer to indicate breaking or iterating that level of enclosing loop; thus

```
break 2
```

exits from two levels of enclosing loops, and **break 1** is equivalent to **break**. **next 2** iterates the second enclosing loop. (Realistically, multi-level **break**'s and **next**'s are not likely to be much used because they lead to code that is hard to understand and somewhat risky to change.)

The "while" Statement

One of the problems with the Fortran DO statement is that it generally insists upon being done once, regardless of its limits. If a loop begins

```
DO I = 2, 1
```

this will typically be done once with **I** set to 2, even though common sense would suggest that perhaps it shouldn't be. Of course a Ratfor **do** can easily be preceded by a test

```

if (j <= k)
  do i = j, k {
    ---
  }

```

but this has to be a conscious act, and is often overlooked by programmers.

A more serious problem with the DO statement is that it encourages that a program be written in terms of an arithmetic progression with small positive steps, even though that may not be the best way to write it. If code has to be contorted to fit the requirements imposed by the Fortran DO, it is that much harder to write and understand.

To overcome these difficulties, Ratfor provides a **while** statement, which is simply a loop: "while some condition is true, repeat this group

of statements". It has no preconceptions about why one is looping. For example, this routine to compute sin(x) by the Maclaurin series combines two termination criteria.

```

real function sin(x, e)
  # returns sin(x) to accuracy e, by
  # sin(x) = x - x**3/3! + x**5/5! - ...

  sin = x
  term = x

  i = 3
  while (abs(term) > e & i < 100) {
    term = -term * x**2 / float(i*(i-1))
    sin = sin + term
    i = i + 2
  }

  return
end

```

Notice that if the routine is entered with **term** already smaller than **e**, the loop will be done *zero times*, that is, no attempt will be made to compute **x**3** and thus a potential underflow is avoided. Since the test is made at the top of a **while** loop instead of the bottom, a special case disappears — the code works at one of its boundaries. (The test **i < 100** is the other boundary — making sure the routine stops after some maximum number of iterations.)

As an aside, a sharp character "#" in a line marks the beginning of a comment, the rest of the line is comment. Comments and code can co-exist on the same line — one can make marginal remarks, which is not possible with Fortran's "C in column 1" convention. Blank lines are also permitted anywhere (they are not in Fortran); they should be used to emphasize the natural divisions of a program.

The syntax of the **while** statement is

```

while (legal Fortran condition)
  Rattor statement

```

As with the **if**, *legal Fortran condition* is something that can go into a Fortran Logical IF, and *Rattor statement* is a single statement, which may be multiple statements in braces.

The **while** encourages a style of coding not normally practiced by Fortran programmers. For example, suppose **nextch** is a function which returns the next input character both as a function value and in its argument. Then a loop to find the first non-blank character is just

```
while (nextch(ich) == iblank)
```

A semicolon by itself is a null statement, which is necessary here to mark the end of the **while**; if it were not present, the **while** would control the next statement. When the loop is broken, **ich** contains the first non-blank. Of course the same code can be written in Fortran as

```
100 if (nextch(ich) .eq. iblank) goto 100
```

but many Fortran programmers (and a few compilers) believe this line is illegal. The language at one's disposal strongly influences how one thinks about a problem.

The "for" Statement

The **for** statement is another Ratfor loop, which attempts to carry the separation of loop-body from reason-for-looping a step further than the **while**. A **for** statement allows explicit initialization and increment steps as part of the statement. For example, a **DO** loop is just

```
for (i = 1; i <= n; i = i + 1) ...
```

This is equivalent to

```
i = 1
while (i <= n) {
    ...
    i = i + 1
}
```

The initialization and increment of **i** have been moved into the **for** statement, making it easier to see at a glance what controls the loop.

The **for** and **while** versions have the advantage that they will be done zero times if **n** is less than 1; this is not true of the **do**.

The loop of the sine routine in the previous section can be re-written with a **for** as

```
for (i=3; abs(term) > e & i < 100; i=i+2) {
    term = -term * x**2 / float(i*(i-1))
    sin = sin + term
}
```

The syntax of the **for** statement is

```
for ( init ; condition ; increment )
    Ratfor statement
```

init is any single Fortran statement, which gets done once before the loop begins. *increment* is any single Fortran statement, which gets done at the end of each pass through the loop, before the test. *condition* is again anything that is legal in a logical IF. Any of *init*, *condition*, and *increment* may be omitted, although the semicolons

must always be present. A non-existent *condition* is treated as always true, so **for(;;)** is an indefinite repeat. (But see the **repeat-until** in the next section.)

The **for** statement is particularly useful for backward loops, chaining along lists, loops that might be done zero times, and similar things which are hard to express with a **DO** statement, and obscure to write out with **IF**'s and **GOTO**'s. For example, here is a backwards **DO** loop to find the last non-blank character on a card:

```
for (i = 80; i > 0; i = i - 1)
    if (card(i) != blank)
        break
```

("!=" is the same as ".NE."). The code scans the columns from 80 through to 1. If a non-blank is found, the loop is immediately broken. (**break** and **next** work in **for**'s and **while**'s just as in **do**'s). If **i** reaches zero, the card is all blank.

This code is rather nasty to write with a regular Fortran **DO**, since the loop must go forward, and we must explicitly set up proper conditions when we fall out of the loop. (Forgetting this is a common error.) Thus:

```
DO 10 J = 1, 80
    I = 81 - J
    IF (CARD(I) .NE. BLANK) GO TO 11
10 CONTINUE
I = 0
11 ...
```

The version that uses the **for** handles the termination condition properly for free; **i** is zero when we fall out of the **for** loop.

The increment in a **for** need not be an arithmetic progression; the following program walks along a list (stored in an integer array **ptr**) until a zero pointer is found, adding up elements from a parallel array of values:

```
sum = 0.0
for (i = first; i > 0; i = ptr(i))
    sum = sum + value(i)
```

Notice that the code works correctly if the list is empty. Again, placing the test at the top of a loop instead of the bottom eliminates a potential boundary error.

The "repeat-until" statement

In spite of the dire warnings, there are times when one really needs a loop that tests at the bottom after one pass through. This service is provided by the **repeat-until**:

```
repeat
  Ratfor statement
until (legal Fortran condition)
```

The *Ratfor statement* part is done once, then the condition is evaluated. If it is true, the loop is exited; if it is false, another pass is made.

The **until** part is optional, so a bare **repeat** is the cleanest way to specify an infinite loop. Of course such a loop must ultimately be broken by some transfer of control such as **stop**, **return**, or **break**, or an implicit stop such as running out of input with a **READ** statement.

As a matter of observed fact[8], the **repeat-until** statement is *much* less used than the other looping constructions; in particular, it is typically outnumbered ten to one by **for** and **while**. Be cautious about using it, for loops that test only at the bottom often don't handle null cases well.

More on break and next

break exits immediately from **do**, **while**, **for**, and **repeat-until**. **next** goes to the test part of **do**, **while** and **repeat-until**, and to the increment step of a **for**.

"return" Statement

The standard Fortran mechanism for returning a value from a function uses the name of the function as a variable which can be assigned to; the last value stored in it is the function value upon return. For example, here is a routine **equal** which returns 1 if two arrays are identical, and zero if they differ. The array ends are marked by the special value **-1**.

```
# equal _ compare str1 to str2;
# return 1 if equal, 0 if not
integer function equal(str1, str2)
integer str1(100), str2(100)
integer i

for (i = 1; str1(i) == str2(i); i = i + 1)
  if (str1(i) == -1) {
    equal = 1
    return
  }
equal = 0
return
end
```

In many languages (e.g., PL/1) one instead says

```
return (expression)
```

to return a value from a function. Since this is

often clearer, Ratfor provides such a **return** statement — in a function **F**, **return(expression)** is equivalent to

```
{ F = expression; return }
```

For example, here is **equal** again:

```
# equal _ compare str1 to str2;
# return 1 if equal, 0 if not
integer function equal(str1, str2)
integer str1(100), str2(100)
integer i

for (i = 1; str1(i) == str2(i); i = i + 1)
  if (str1(i) == -1)
    return(1)
return(0)
end
```

If there is no parenthesized expression after **return**, a normal **RETURN** is made. (Another version of **equal** is presented shortly.)

Cosmetics

As we said above, the visual appearance of a language has a substantial effect on how easy it is to read and understand programs. Accordingly, Ratfor provides a number of cosmetic facilities which may be used to make programs more readable.

Free-form Input

Statements can be placed anywhere on a line; long statements are continued automatically, as are long conditions in **if**, **while**, **for**, and **until**. Blank lines are ignored. Multiple statements may appear on one line, if they are separated by semicolons. No semicolon is needed at the end of a line, if Ratfor can make some reasonable guess about whether the statement ends there. Lines ending with any of the characters

```
= + - * , | & ( _
```

are assumed to be continued on the next line. Underscores are discarded wherever they occur; all others remain as part of the statement.

Any statement that begins with an all-numeric field is assumed to be a Fortran label, and placed in columns 1-5 upon output. Thus

```
write(6, 100); 100 format("hello")
```

is converted into

```

write(6, 100)
100 format(5hhello)

```

Translation Services

Text enclosed in matching single or double quotes is converted to `nH...` but is otherwise unaltered (except for formatting — it may get split across card boundaries during the reformatting process). Within quoted strings, the backslash `\` serves as an escape character: the next character is taken literally. This provides a way to get quotes (and of course the backslash itself) into quoted strings:

```
"\\"
```

is a string containing a backslash and an apostrophe. (This is *not* the standard convention of doubled quotes, but it is easier to use and more general.)

Any line that begins with the character `%` is left absolutely unaltered except for stripping off the `%` and moving the line one position to the left. This is useful for inserting control cards, and other things that should not be transmogrified (like an existing Fortran program). Use `%` only for ordinary statements, not for the condition parts of `if`, `while`, etc., or the output may come out in an unexpected place.

The following character translations are made, except within single or double quotes or on a line beginning with a `%`:

<code>==</code>	<code>.eq.</code>	<code>!=</code>	<code>.ne.</code>
<code>></code>	<code>.gt.</code>	<code>>=</code>	<code>.ge.</code>
<code><</code>	<code>.lt.</code>	<code><=</code>	<code>.le.</code>
<code>&</code>	<code>.and.</code>	<code> </code>	<code>.or.</code>
<code>!</code>	<code>.not.</code>	<code>~</code>	<code>.not.</code>

In addition, the following translations are provided for input devices with restricted character sets.

<code>{</code>	<code>{</code>	<code>}</code>	<code>}</code>
<code>\$</code>	<code>{</code>	<code>\$</code>	<code>}</code>

"define" Statement

Any string of alphanumeric characters can be defined as a name; thereafter, whenever that name occurs in the input (delimited by non-alphanumerics) it is replaced by the rest of the definition line. (Comments and trailing white spaces are stripped off). A defined name can be arbitrarily long, and must begin with a letter.

`define` is typically used to create symbolic parameters:

```

define ROWS 100
define COLS 50
dimension a(ROWS), b(ROWS, COLS)
      if (i > ROWS | j > COLS) ...

```

Alternately, definitions may be written as

```
define(ROWS, 100)
```

In this case, the defining text is everything after the comma up to the balancing right parenthesis; this allows multi-line definitions.

It is generally a wise practice to use symbolic parameters for most constants, to help make clear the function of what would otherwise be mysterious numbers. As an example, here is the routine `equal` again, this time with symbolic constants.

```

define YES 1
define NO 0
define EOS -1
define ARB 100

```

```

# equal _ compare str1 to str2;
# return YES if equal, NO if not
integer function equal(str1, str2)
integer str1(ARB), str2(ARB)
integer i

```

```

for (i = 1; str1(i) == str2(i); i = i + 1)
  if (str1(i) == EOS)
    return(YES)
return(NO)
end

```

"include" Statement

The statement

```
include file
```

inserts the file found on input stream *file* into the Ratfor input in place of the `include` statement. The standard usage is to place `COMMON` blocks on a file, and `include` that file whenever a copy is needed:

```

subroutine x
  include commonblocks
  ...
end

suroutine y
  include commonblocks
  ...
end

```

This ensures that all copies of the `COMMON` blocks are identical

Pitfalls, Botches, Blemishes and other Fallings

Ratfor catches certain syntax errors, such as missing braces, **else** clauses without an **if**, and most errors involving missing parentheses in statements. Beyond that, since Ratfor knows no Fortran, any errors you make will be reported by the Fortran compiler, so you will from time to time have to relate a Fortran diagnostic back to the Ratfor source.

Keywords are reserved — using **if**, **else**, etc., as variable names will typically wreak havoc. Don't leave spaces in keywords. Don't use the Arithmetic IF.

The Fortran **nH** convention is not recognized anywhere by Ratfor; use quotes instead.

3. IMPLEMENTATION

Ratfor was originally written in C[4] on the UNIX operating system[5]. The language is specified by a context free grammar and the compiler constructed using the YACC compiler-compiler[6].

The Ratfor grammar is simple and straightforward, being essentially

```

prog  : stat
      | prog stat
stat  : if (...) stat
      | if (...) stat else stat
      | while (...) stat
      | for (...) stat
      | do ... stat
      | repeat stat
      | repeat stat until (...)
      | switch (...) { case ...: prog ...
                      default: prog }
      | return
      | break
      | next
      | digits stat
      | { prog }
      | anything unrecognizable

```

The observation that Ratfor knows no Fortran follows directly from the rule that says a statement is "anything unrecognizable". In fact most of Fortran falls into this category, since any statement that does not begin with one of the keywords is by definition "unrecognizable."

Code generation is also simple. If the first thing on a source line is not a keyword (like **if**, **else**, etc.) the entire statement is simply copied to the output with appropriate character translation and formatting. (Leading digits are treated as a label.) Keywords cause only slightly more complicated actions. For example, when **if** is

recognized, two consecutive labels **L** and **L+1** are generated and the value of **L** is stacked. The condition is then isolated, and the code

```
if (.not. (condition)) goto L
```

is output. The *statement* part of the **if** is then translated. When the end of the statement is encountered (which may be some distance away and include nested **if**'s, of course), the code

```
L    continue
```

is generated, unless there is an **else** clause, in which case the code is

```
      goto L+1
L    continue
```

In this latter case, the code

```
L+1  continue
```

is produced after the *statement* part of the **else**. Code generation for the various loops is equally simple.

One might argue that more care should be taken in code generation. For example, if there is no trailing **else**,

```
if (i > 0) x = a
```

should be left alone, not converted into

```
if (.not. (i .gt. 0)) goto 100
x = a
100  continue
```

But what are optimizing compilers for, if not to improve code? It is a rare program indeed where this kind of "inefficiency" will make even a measurable difference. In the few cases where it is important, the offending lines can be protected by '%'.
The use of a compiler-compiler is definitely the preferred method of software development. The language is well-defined, with few syntactic irregularities. Implementation is quite simple; the original construction took under a week. The language is sufficiently simple, however, that an *ad hoc* recognizer can be readily constructed to do the same job if no compiler-compiler is available.

The C version of Ratfor is used on UNIX and on the Honeywell GCOS systems. C compilers are not as widely available as Fortran, however, so there is also a Ratfor written in itself and originally bootstrapped with the C version. The Ratfor version was written so as to translate into the portable subset of Fortran described in [1], so it is portable, having been run essentially without change on at least twelve

distinct machines. (The main restrictions of the portable subset are: only one character per machine word; subscripts in the form $c \pm v \pm c$; avoiding expressions in places like DO loops; consistency in subroutine argument usage, and in COMMON declarations. Ratfor itself will not gratuitously generate non-standard Fortran.)

The Ratfor version is about 1500 lines of Ratfor (compared to about 1000 lines of C); this compiles into 2500 lines of Fortran. This expansion ratio is somewhat higher than average, since the compiled code contains unnecessary occurrences of COMMON declarations. The execution time of the Ratfor version is dominated by two routines that read and write cards. Clearly these routines could be replaced by machine coded local versions; unless this is done, the efficiency of other parts of the translation process is largely irrelevant.

4. EXPERIENCE

Good Things

"It's so much better than Fortran" is the most common response of users when asked how well Ratfor meets their needs. Although cynics might consider this to be vacuous, it does seem to be true that decent control flow and cosmetics converts Fortran from a bad language into quite a reasonable one, assuming that Fortran data structures are adequate for the task at hand.

Although there are no quantitative results, users feel that coding in Ratfor is at least twice as fast as in Fortran. More important, debugging and subsequent revision are much faster than in Fortran. Partly this is simply because the code can be read. The looping statements which test at the top instead of the bottom seem to eliminate or at least reduce the occurrence of a wide class of boundary errors. And of course it is easy to do structured programming in Ratfor; this self-discipline also contributes markedly to reliability.

One interesting and encouraging fact is that programs written in Ratfor tend to be as

readable as programs written in more modern languages like Pascal. Once one is freed from the shackles of Fortran's clerical detail and rigid input format, it is easy to write code that is readable, even esthetically pleasing. For example, here is a Ratfor implementation of the linear table search discussed by Knuth [7]:

```
A(m+1) = x
for (i = 1; A(i) != x; i = i + 1)
;
if (i > m) {
    m = i
    B(i) = 1
}
else
    B(i) = B(i) + 1
```

A large corpus (5400 lines) of Ratfor, including a subset of the Ratfor preprocessor itself, can be found in [8].

Bad Things

The biggest single problem is that many Fortran syntax errors are not detected by Ratfor but by the local Fortran compiler. The compiler then prints a message in terms of the generated Fortran, and in a few cases this may be difficult to relate back to the offending Ratfor line, especially if the implementation conceals the generated Fortran. This problem could be dealt with by tagging each generated line with some indication of the source line that created it, but this is inherently implementation-dependent, so no action has yet been taken. Error message interpretation is actually not so arduous as might be thought. Since Ratfor generates no variables, only a simple pattern of IF's and GOTO's, data-related errors like missing DIMENSION statements are easy to find in the Fortran. Furthermore, there has been a steady improvement in Ratfor's ability to catch trivial syntactic errors like unbalanced parentheses and quotes.

There are a number of implementation weaknesses that are a nuisance, especially to new users. For example, keywords are reserved. This rarely makes any difference, except for those hardy souls who want to use an Arithmetic IF. A few standard Fortran constructions are not accepted by Ratfor, and this is perceived as a problem by users with a large corpus of existing Fortran programs. Protecting every line with a '%' is not really a complete solution, although it serves as a stop-gap. The best long-term solution is provided by the program Struct [9], which converts arbitrary Fortran programs into Ratfor.

Users who export programs often complain

that the generated Fortran is "unreadable" because it is not tastefully formatted and contains extraneous CONTINUE statements. To some extent this can be ameliorated (Ratfor now has an option to copy Ratfor comments into the generated Fortran), but it has always seemed that effort is better spent on the input language than on the output esthetics.

One final problem is partly attributable to success — since Ratfor is relatively easy to modify, there are now several dialects of Ratfor. Fortunately, so far most of the differences are in character set, or in invisible aspects like code generation.

5. CONCLUSIONS

Ratfor demonstrates that with modest effort it is possible to convert Fortran from a bad language into quite a good one. A preprocessor is clearly a useful way to extend or ameliorate the facilities of a base language.

When designing a language, it is important to concentrate on the essential requirement of providing the user with the best language possible for a given effort. One must avoid throwing in "features" — things which the user may trivially construct within the existing framework.

One must also avoid getting sidetracked on irrelevancies. For instance it seems pointless for Ratfor to prepare a neatly formatted listing of either its input or its output. The user is presumably capable of the self-discipline required to prepare neat input that reflects his thoughts. It is much more important that the language provide free-form input so he *can* format it neatly. No one should read the output anyway except in the most dire circumstances.

Acknowledgements

C. A. R. Hoare once said that "One thing [the language designer] should not do is to include untried ideas of his own." Ratfor follows this precept very closely — everything in it has been stolen from someone else. Most of the control flow structures are taken directly from the language C[4] developed by Dennis Ritchie; the comment and continuation conventions are adapted from Altran[10].

I am grateful to Stuart Feldman, whose patient simulation of an innocent user during the early days of Ratfor led to several design improvements and the eradication of bugs. He also translated the C parse-tables and YACC parser into Fortran for the first Ratfor version of Ratfor.

References

- [1] B. G. Ryder, "The PFORT Verifier," *Software—Practice & Experience*, October 1974.
- [2] American National Standard Fortran. American National Standards Institute, New York, 1966.
- [3] *For-word: Fortran Development Newsletter*, August 1975.
- [4] D. M. Ritchie, B. W. Kernighan and M. E. Lesk, "The C Programming Language," Bell Laboratories Computing Science Technical Report #31, 1975.
- [5] D. M. Ritchie and K. L. Thompson, "The UNIX Time-sharing System." *CACM*, July 1974.
- [6] S. C. Johnson, "YACC — Yet Another Compiler-Compiler." Bell Laboratories Computing Science Technical Report #32, 1974.
- [7] D. E. Knuth, "Structured Programming with goto Statements." *Computing Surveys*, December 1974.
- [8] B. W. Kernighan and P. J. Plauger, *Software Tools*, Addison-Wesley, 1976.
- [9] B. S. Baker, "Struct — A Program which Structures Fortran", Bell Laboratories internal memorandum, December 1975.
- [10] A. D. Hall, "The Altran System for Rational Function Manipulation — A Survey." *CACM*, August 1971.

Appendix: Usage on UNIX and GCOS.

Beware — local customs vary. Check with a native before going into the jungle.

UNIX

The program **ratfor** is the basic translator; it takes either a list of file names or the standard input and writes Fortran on the standard output. Options include **-6x**, which uses **x** as a continuation character in column 6 (UNIX uses **&** in column 1), and **-C**, which causes Fortran comments to be copied into the generated Fortran.

The program **rc** provides an interface to the **ratfor** command which is much the same as **cc**. Thus

```
rc [options] files
```

compiles the files specified by **files**. Files with names ending in **.r** are Ratfor source; other files are assumed to be for the loader. The flags **-C** and **-6x** described above are recognized, as are

- c** compile only; don't load
- f** save intermediate Fortran **.f** files
- r** Ratfor only; implies **-c** and **-f**
- 2** use big Fortran compiler (for large programs)
- U** flag undeclared variables (not universally available)

Other flags are passed on to the loader.

GCOS

The program **./ratfor** is the bare translator, and is identical to the UNIX version, except that the continuation convention is **&** in column 6. Thus

```
./ratfor files >output
```

translates the Ratfor source on **files** and collects the generated Fortran on file 'output' for subsequent processing.


./rc provides much the same services as **rc** (within the limitations of GCOS), regrettably with a somewhat different syntax. Options recognized by **./rc** include


name	Ratfor source or library, depending on type
h=/name	make TSS H+ file (runnable version); run as /name
r=/name	update and use random library
a=	compile as ascii (default is bcd)
C=	copy comments into Fortran
f=name	Fortran source file
g=name	gmap source file

Other options are as specified for the **./cc** command described in [4].

TSO, TSS, and other systems

Ratfor exists on various other systems; check with the author for specifics.





The M4 Macro Processor

Brian W. Kernighan

Dennis M. Ritchie

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

M4 is a macro processor available on UNIX and GCOS. Its primary use has been as a front end for Ratfor for those cases where parameterless macros are not adequately powerful. It has also been used for languages as disparate as C and Cobol. M4 is particularly suited for functional languages like Fortran, PL/I and C since macros are specified in a functional notation.

M4 provides features seldom found even in much larger macro processors, including

- arguments
- condition testing
- arithmetic capabilities
- string and substring functions
- file manipulation

This paper is a user's manual for M4.

April 1, 1977



The M4 Macro Processor

Brian W. Kernighan

Dennis M. Ritchie

Bell Laboratories
Murray Hill, New Jersey 07974

Introduction

A macro processor is a useful way to enhance a programming language, to make it more palatable or more readable, or to tailor it to a particular application. The **#define** statement in C and the analogous **define** in Ratfor are examples of the basic facility provided by any macro processor — replacement of text by other text.

The M4 macro processor is an extension of a macro processor called M3 which was written by D. M. Ritchie for the AP-3 minicomputer; M3 was in turn based on a macro processor implemented for [1]. Readers unfamiliar with the basic ideas of macro processing may wish to read some of the discussion there.

M4 is a suitable front end for Ratfor and C, and has also been used successfully with Cobol. Besides the straightforward replacement of one string of text by another, it provides macros with arguments, conditional macro expansion, arithmetic, file manipulation, and some specialized string processing functions.

The basic operation of M4 is to copy its input to its output. As the input is read, however, each alphanumeric “token” (that is, string of letters and digits) is checked. If it is the name of a macro, then the name of the macro is replaced by its defining text, and the resulting string is pushed back onto the input to be rescanned. Macros may be called with arguments, in which case the arguments are collected and substituted into the right places in the defining text before it is rescanned.

M4 provides a collection of about twenty built-in macros which perform various useful operations; in addition, the user can define new macros. Built-ins and user-defined macros work exactly the same way,

except that some of the built-in macros have side effects on the state of the process.

Usage

On UNIX, use

m4 [files]

Each argument file is processed in order; if there are no arguments, or if an argument is ‘-’, the standard input is read at that point. The processed text is written on the standard output, which may be captured for subsequent processing with

m4 [files] >outputfile

On GCOS, usage is identical, but the program is called **./m4**.

Defining Macros

The primary built-in function of M4 is **define**, which is used to define new macros. The input

define(name, stuff)

causes the string **name** to be defined as **stuff**. All subsequent occurrences of **name** will be replaced by **stuff**. **name** must be alphanumeric and must begin with a letter (the underscore **_** counts as a letter). **stuff** is any text that contains balanced parentheses; it may stretch over multiple lines.

Thus, as a typical example,

define(N, 100)

...

if (i > N)

defines **N** to be 100, and uses this “symbolic constant” in a later **if** statement.

The left parenthesis must immediately follow the word **define**, to signal that **define** has arguments. If a macro or built-in name

is not followed immediately by `(`, it is assumed to have no arguments. This is the situation for **N** above; it is actually a macro **with no arguments**, and thus when it is used there need be no (...) following it.

You should also notice that a macro name is only recognized as such if it appears surrounded by non-alphanumerics. For example, in

```
define(N, 100)
...
if (NNN > 100)
```

the variable **NNN** is absolutely unrelated to the defined macro **N**, even though it contains a lot of **N**'s.

Things may be defined in terms of other things. For example,

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

defines both **M** and **N** to be 100.

What happens if **N** is redefined? Or, to say it another way, is **M** defined as **N** or as 100? In M4, the latter is true — **M** is 100, so even if **N** subsequently changes, **M** does not.

This behavior arises because M4 expands macro names into their defining text as soon as it possibly can. Here, that means that when the string **N** is seen as the arguments of **define** are being collected, it is immediately replaced by 100; it's just as if you had said

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

in the first place.

If this isn't what you really want, there are two ways out of it. The first, which is specific to this situation, is to interchange the order of the definitions:

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

Now **M** is defined to be the string **N**, so when you ask for **M** later, you'll always get the value of **N** at that time (because the **M** will be replaced by **N** which will be replaced by 100).

Quoting

The more general solution is to delay the expansion of the arguments of **define** by quoting them. Any text surrounded by the single quotes ` and ` is not expanded immediately, but has the quotes stripped off. If you say

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

the quotes around the **N** are stripped off as the argument is being collected, but they have served their purpose, and **M** is defined as the string **N**, not 100. The general rule is that M4 always strips off one level of single quotes whenever it evaluates something. This is true even outside of macros. If you want the word **define** to appear in the output, you have to quote it in the input, as in

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

As another instance of the same thing, which is a bit more surprising, consider redefining **N**:

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

Perhaps regrettably, the **N** in the second definition is evaluated as soon as it's seen; that is, it is replaced by 100, so it's as if you had written

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

This statement is ignored by M4, since you can only define things that look like names, but it obviously doesn't have the effect you wanted. To really redefine **N**, you must delay the evaluation by quoting:

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

In M4, it is often wise to quote the first argument of a macro.

If ` and ` are not convenient for some reason, the quote characters can be changed with the built-in **changequote**:

```
changequote(l, l)
```

makes the new quote characters the left and right brackets. You can restore the original characters with just

changequote

There are two additional built-ins related to **define**. **undefine** removes the definition of some macro or built-in:

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

removes the definition of N. (Why are the quotes absolutely necessary?) Built-ins can be removed with **undefine**, as in

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

but once you remove one, you can never get it back.

The built-in **ifdef** provides a way to determine if a macro is currently defined. In particular, M4 has pre-defined the names **unix** and **gcos** on the corresponding systems, so you can tell which one you're using:

```
ifdef('unix', 'define(wordsize,16)')
ifdef('gcos', 'define(wordsize,36)')
```

makes a definition appropriate for the particular machine. Don't forget the quotes!

ifdef actually permits three arguments; if the name is undefined, the value of **ifdef** is then the third argument, as in

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

Arguments

So far we have discussed the simplest form of macro processing — replacing one string by another (fixed) string. User-defined macros may also have arguments, so different invocations can have different results. Within the replacement text for a macro (the second argument of its **define**) any occurrence of **\$n** will be replaced by the **n**th argument when the macro is actually used. Thus, the macro **bump**, defined as

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

generates code to increment its argument by 1:

```
bump(x)
```

is

```
x = x + 1
```

A macro can have as many arguments as you want, but only the first nine are accessible, through **\$1** to **\$9**. (The macro

name itself is **\$0**, although that is less commonly used.) Arguments that are not supplied are replaced by null strings, so we can define a macro **cat** which simply concatenates its arguments, like this:

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

Thus

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

is equivalent to

```
xyz
```

\$4 through **\$9** are null, since no corresponding arguments were provided.

Leading unquoted blanks, tabs, or newlines that occur during argument collection are discarded. All other white space is retained. Thus

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

defines **a** to be **b c**.

Arguments are separated by commas, but parentheses are counted properly, so a comma "protected" by parentheses does not terminate an argument. That is, in

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

there are only two arguments; the second is literally **(b,c)**. And of course a bare comma or parenthesis can be inserted by quoting it.

Arithmetic Built-ins

M4 provides two built-in functions for doing arithmetic on integers (only). The simplest is **incr**, which increments its numeric argument by 1. Thus to handle the common programming situation where you want a variable to be defined as "one more than N", write

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

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

The more general mechanism for arithmetic is a built-in called **eval**, which is capable of arbitrary arithmetic on integers. It provides the operators (in decreasing order of precedence)

unary + and
 ** or ^ (exponentiation)
 * / % (modulus)
 + -
 == != < <= > >=
 ! (not)
 & or && (logical and)
 | or || (logical or)

Parentheses may be used to group operations where needed. All the operands of an expression given to **eval** must ultimately be numeric. The numeric value of a true relation (like $1 > 0$) is 1, and false is 0. The precision in **eval** is 32 bits on UNIX and 36 bits on GCOS.

As a simple example, suppose we want **M** to be $2**N+1$. Then

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

As a matter of principle, it is advisable to quote the defining text for a macro unless it is very simple indeed (say just a number); it usually gives the result you want, and is a good habit to get into.

File Manipulation

You can include a new file in the input at any time by the built-in function **include**:

```
include(filename)
```

inserts the contents of **filename** in place of the **include** command. The contents of the file is often a set of definitions. The value of **include** (that is, its replacement text) is the contents of the file; this can be captured in definitions, etc.

It is a fatal error if the file named in **include** cannot be accessed. To get some control over this situation, the alternate form **sinclude** can be used; **sinclude** ("silent include") says nothing and continues if it can't access the file.

It is also possible to divert the output of M4 to temporary files during processing, and output the collected material upon command. M4 maintains nine of these diversions, numbered 1 through 9. If you say

```
divert(n)
```

all subsequent output is put onto the end of a temporary file referred to as **n**. Diverting to this file is stopped by another **divert** com-

mand; in particular, **divert** or **divert(0)** resumes the normal output process.

Diverted text is normally output all at once at the end of processing, with the diversions output in numeric order. It is possible, however, to bring back diversions at any time, that is, to append them to the current diversion.

undivert

brings back all diversions in numeric order, and **undivert** with arguments brings back the selected diversions in the order given. The act of undiverting discards the diverted stuff, as does diverting into a diversion whose number is not between 0 and 9 inclusive.

The value of **undivert** is *not* the diverted stuff. Furthermore, the diverted material is *not* rescanned for macros.

The built-in **divnum** returns the number of the currently active diversion. This is zero during normal processing.

System Command

You can run any program in the local operating system with the **syscmd** built-in. For example,

```
syscmd(date)
```

on UNIX runs the **date** command. Normally **syscmd** would be used to create a file for a subsequent **include**.

To facilitate making unique file names, the built-in **maketemp** is provided, with specifications identical to the system function *mktemp*: a string of XXXXX in the argument is replaced by the process id of the current process.

Conditionals

There is a built-in called **ifelse** which enables you to perform arbitrary conditional testing. In the simplest form,

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

compares the two strings **a** and **b**. If these are identical, **ifelse** returns the string **c**; otherwise it returns **d**. Thus we might define a macro called **compare** which compares two strings and returns "yes" or "no" if they are the same or different.

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

Note the quotes, which prevent too-early evaluation of **ifelse**.

If the fourth argument is missing, it is treated as empty.

ifelse can actually have any number of arguments, and thus provides a limited form of multi-way decision capability. In the input

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

if the string **a** matches the string **b**, the result is **c**. Otherwise, if **d** is the same as **e**, the result is **f**. Otherwise the result is **g**. If the final argument is omitted, the result is null, so

ifelse(a, b, c)

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

String Manipulation

The built-in **len** returns the length of the string that makes up its argument. Thus

len(abcdef)

is 6, and **len((a,b))** is 5.

The built-in **substr** can be used to produce substrings of strings. **substr(s, i, n)** returns the substring of **s** that starts at the **i**th position (origin zero), and is **n** characters long. If **n** is omitted, the rest of the string is returned, so

substr('now is the time', 1)

is

ow is the time

If **i** or **n** are out of range, various sensible things happen.

index(s1, s2) returns the index (position) in **s1** where the string **s2** occurs, or -1 if it doesn't occur. As with **substr**, the origin for strings is 0.

The built-in **translit** performs character transliteration.

translit(s, f, t)

modifies **s** by replacing any character found in **f** by the corresponding character of **t**. That is,

translit(s, aeiou, 12345)

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

translit(s, aeiou)

deletes vowels from **s**.

There is also a built-in called **dnl** which deletes all characters that follow it up to and including the next newline; it is useful mainly for throwing away empty lines that otherwise tend to clutter up M4 output. For example, if you say

define(N, 100)

define(M, 200)

define(L, 300)

the newline at the end of each line is not part of the definition, so it is copied into the output, where it may not be wanted. If you add **dnl** to each of these lines, the newlines will disappear.

Another way to achieve this, due to J. E. Weythman, is

divert(-1)

define(...)

...

divert

Printing

The built-in **errprint** writes its arguments out on the standard error file. Thus you can say

errprint('fatal error')

dumpdef is a debugging aid which dumps the current definitions of defined terms. If there are no arguments, you get everything; otherwise you get the ones you name as arguments. Don't forget to quote the names!

Summary of Built-ins

Each entry is preceded by the page number where it is described.

```
3  changequote(L, R)
1  define(name, replacement)
4  divert(number)
4  divnum
5  dnl
5  dumpdef('name', 'name', ...)
5  errprint(s, s, ...)
4  eval(numeric expression)
3  ifdef('name', this if true, this if false)
5  ifelse(a, b, c, d)
4  include(file)
3  incr(number)
5  index(s1, s2)
5  len(string)
4  maketemp(...XXXXX...)
4  sinclude(file)
5  substr(string, position, number)
4  syscmd(s)
5  translit(str, from, to)
3  undefine('name')
4  undivert(number,number,...)
```


Acknowledgements

We are indebted to Rick Becker, John Chambers, Doug McIlroy, and Jim Weythman, whose pioneering use of M4 has led to several valuable improvements. We are also deeply grateful to Weythman for several substantial contributions to the code.

References

- [1] B. W. Kernighan and P. J. Plauger, *Software Tools*, Addison-Wesley, Inc., 1976.





Make — A Program for Maintaining Computer Programs

S. I. Feldman

Bell Laboratories
Murray Hill, New Jersey 07974

Introduction

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

The program described in this report mechanizes many of the activities of program development and maintenance. If the information on inter-file dependences and command sequences is stored in a file, the simple command

```
make
```

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

```
think — edit — make — test . . .
```

Make is most useful for medium-sized programming projects; it does not solve the problems of maintaining multiple source versions or of describing huge programs. *Make* was designed for use on Unix, but a version runs on GCOS.

Basic Features

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

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

```
#include "defs"
```

The following text describes the relationships and operations:

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

x.o y.o : defs
```

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

```
make
```

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

Make operates using three sources of information: a user-supplied description file (as above), file names and "last-modified" times from the file system, and built-in rules to bridge some of the gaps. In our example, the first line says that *prog* depends on three ".o" files. Once these object files are current, the second line describes how to load them to create *prog*. The third line says that *x.o* and *y.o* depend on the file *defs*. From the file system, *make* discovers that there are three ".c" files corresponding to the needed ".o" files, and uses built-in information on how to generate an object from a source file (*i.e.*, issue a "cc -c" command).

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

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

x.o : x.c defs
      cc -c x.c

y.o : y.c defs
      cc -c y.c

z.o : z.c
      cc -c z.c
```

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

```
make
```

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

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

```
make x.o
```

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

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

Make has a simple macro mechanism for substituting in dependency lines and command

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

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

The last two invocations are identical. \$\$ is a dollar sign. All of these macros are assigned values during input, as shown below. Four special macros change values during the execution of the command: \$*, \$@, \$?, and \$<. They will be discussed later. The following fragment shows the use:

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

The command

```
make
```

loads the three object files with the *lS* library. The command

```
make "LIBES = -ll -lp"
```

loads them with both the Lex (“-ll”) and the portable (“-lp”) libraries, since macro definitions on the command line override definitions in the description. (It is necessary to quote arguments with embedded blanks in Unix commands.)

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

Description Files and Substitutions

A description file contains three types of information: macro definitions, dependency information, and executable commands. There is also a comment convention: all characters after a sharp (#) are ignored, as in the sharp itself. Blank lines and lines beginning with a sharp are totally ignored. If a non-comment line is too long, it can be continued using a backslash. If the last character of a line is a backslash, the backslash, newline, and following blanks and tabs are replaced by a single blank.

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

```
2 = xyz
abc = -ll -ly -lp
LIBES =
```

The last definition assigns LIBES the null string. A macro that is never explicitly defined has the null string as value. Macro definitions may also appear on the *make* command line (see below).

Other lines give information about target files. The general form of an entry is:

```
target1 [target2 . . .] [:] [dependent1 . . .] [; commands] [# . . .]  
[(tab) commands] [# . . .]  
...
```

Items inside brackets may be omitted. Targets and dependents are strings of letters, digits, periods, and slashes. (Shell metacharacters "*" and "?" are expanded.) A command is any string of characters not including a sharp (except in quotes) or newline. Commands may appear either after a semicolon on a dependency line or on lines beginning with a tab immediately following a dependency line.

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

1. For the usual single-colon case, at most one of these dependency lines may have a command sequence associated with it. If the target is out of date with any of the dependents on any of the lines, and a command sequence is specified (even a null one following a semicolon or tab), it is executed; otherwise a default creation rule may be invoked.
2. In the double-colon case, a command sequence may be associated with each dependency line; if the target is out of date with any of the files on a particular line, the associated commands are executed. A built-in rule may also be executed. This detailed form is of particular value in updating archive-type files.

If a target must be created, the sequence of commands is executed. Normally, each command line is printed and then passed to a separate invocation of the Shell after substituting for macros. (The printing is suppressed in silent mode or if the command line begins with an @ sign). *Make* normally stops if any command signals an error by returning a non-zero error code. (Errors are ignored if the "-i" flag has been specified on the *make* command line, if the fake target name ".IGNORE" appears in the description file, or if the command string in the description file begins with a hyphen. Some Unix commands return meaningless status). Because each command line is passed to a separate invocation of the Shell, care must be taken with certain commands (e.g., *chdir* and Shell control commands) that have meaning only within a single Shell process; the results are forgotten before the next line is executed.

Before issuing any command, certain macros are set. \$@ is set to the name of the file to be "made". \$? is set to the string of names that were found to be younger than the target. If the command was generated by an implicit rule (see below), \$< is the name of the related file that caused the action, and \$* is the prefix shared by the current and the dependent file names.

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

Command Usage

The *make* command takes four kinds of arguments: macro definitions, flags, description file names, and target file names.

```
make [ flags ] [ macro definitions ] [ targets ]
```

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

First, all macro definition arguments (arguments with embedded equal signs) are analyzed and the assignments made. Command-line macros override corresponding definitions found in the description files.

Next, the flag arguments are examined. The permissible flags are

- i Ignore error codes returned by invoked commands. This mode is entered if the fake target name ".IGNORE" appears in the description file.
- s Silent mode. Do not print command lines before executing. This mode is also entered if the fake target name ".SILENT" appears in the description file.
- r Do not use the built-in rules.
- n No execute mode. Print commands, but do not execute them. Even lines beginning with an "@" sign are printed.
- t Touch the target files (causing them to be up to date) rather than issue the usual commands.
- q Question. The *make* command returns a zero or non-zero status code depending on whether the target file is or is not up to date.
- p Print out the complete set of macro definitions and target descriptions
- d Debug mode. Print out detailed information on files and times examined.
- f Description file name. The next argument is assumed to be the name of a description file. A file name of "-" denotes the standard input. If there are no "-f" arguments, the file named *makefile* or *Makefile* in the current directory is read. The contents of the description files override the built-in rules if they are present).


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

Implicit Rules

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

<i>.o</i>	Object file
<i>.c</i>	C source file
<i>.e</i>	Efl source file
<i>.r</i>	Ratfor source file
<i>.f</i>	Fortran source file
<i>.s</i>	Assembler source file
<i>.y</i>	Yacc-C source grammar
<i>.yr</i>	Yacc-Ratfor source grammar
<i>.ye</i>	Yacc-Efl source grammar
<i>.l</i>	Lex source grammar

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



If the file *x.o* were needed and there were an *x.c* in the description or directory, it would be compiled. If there were also an *x.l*, that grammar would be run through Lex before

compiling the result. However, if there were no *x.c* but there were an *x.l*, *make* would discard the intermediate C-language file and use the direct link in the graph above.

It is possible to change the names of some of the compilers used in the default, or the flag arguments with which they are invoked by knowing the macro names used. The compiler names are the macros AS, CC, RC, EC, YACC, YACCR, YACCE, and LEX. The command

```
make CC=newcc
```

will cause the "newcc" command to be used instead of the usual C compiler. The macros CFLAGS, RFLAGS, EFLAGS, YFLAGS, and LFLAGS may be set to cause these commands to be issued with optional flags. Thus,

```
make "CFLAGS = -O"
```

causes the optimizing C compiler to be used.

Example

As an example of the use of *make*, we will present the description file used to maintain the *make* command itself. The code for *make* is spread over a number of C source files and a Yacc grammar. The description file contains:

```
# Description file for the Make command
P = und -3 |opr -r2      # send to GCOS to be printed
FILES = Makefile version.c defs main.c doname.c misc.c files.c dosys.cgram.y lex.c gcos.c
OBJECTS = version.o main.o doname.o misc.o files.o dosys.o gram.o
LIBES = -IS
LINT = lint -p
CFLAGS = -O
make: $(OBJECTS)
      cc $(CFLAGS) $(OBJECTS) $(LIBES) -o make
      size make
$(OBJECTS): defs
gram.o: lex.c
cleanup:
      -rm *.o gram.c
      -du
install:
      @size make /usr/bin/make
      cp make /usr/bin/make ; rm make
print: $(FILES)      # print recently changed files
      pr $? |$P
      touch print
test:
      make -dp |grep -v TIME >1zap
      /usr/bin/make -dp |grep -v TIME >2zap
      diff 1zap 2zap
      rm 1zap 2zap
lint : dosys.c doname.c files.c main.c misc.c version.c gram.c
      $(LINT) dosys.c doname.c files.c main.c misc.c version.c gram.c
      rm gram.c
arch:
      ar uv /sys/source/s2/make.a $(FILES)
```

Make usually prints out each command before issuing it. The following output results from

typing the simple command

make

in a directory containing only the source and description file:

```
cc -c version.c
cc -c main.c
cc -c doname.c
cc -c misc.c
cc -c files.c
cc -c dosys.c
yacc gram.y
mv y.tab.c gram.c
cc -c gram.c
cc version.o main.o doname.o misc.o files.o dosys.o gram.o -ls -o make
13188+3348+3044 = 19580b = 046174b
```

Although none of the source files or grammars were mentioned by name in the description file, *make* found them using its suffix rules and issued the needed commands. The string of digits results from the "size make" command; the printing of the command line itself was suppressed by an @ sign. The @ sign on the *size* command in the description file suppressed the printing of the command, so only the sizes are written.

The last few entries in the description file are useful maintenance sequences. The "print" entry prints only the files that have been changed since the last "make print" command. A zero-length file *print* is maintained to keep track of the time of the printing; the \$? macro in the command line then picks up only the names of the files changed since *print* was touched. The printed output can be sent to a different printer or to a file by changing the definition of the *P* macro:

```
make print "P = opr -sp"
      or
make print "P= cat >zap"
```

Suggestions and Warnings

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

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

```
make -n
```

orders *make* to print out the commands it would issue without actually taking the time to execute them. If a change to a file is absolutely certain to be benign (e.g., adding a new definition to an include file), the "-t" (touch) option can save a lot of time: instead of issuing a large number of superfluous recompilations, *make* updates the modification times on the affected file. Thus, the command

```
make -ts
```

("touch silently") causes the relevant files to appear up to date. Obvious care is necessary, since this mode of operation subverts the intention of *make* and destroys all memory of the previous relationships.

The debugging flag ("-d") causes *make* to print out a very detailed description of what it is doing, including the file times. The output is verbose, and recommended only as a last resort.

Acknowledgments

I would like to thank S. C. Johnson for suggesting this approach to program maintenance control. I would like to thank S. C. Johnson and H. Gajewska for being the prime guinea pigs during development of *make*.

References

1. S. J. Johnson, "Yacc — Yet Another Compiler-Compiler", *Computing Science Technical Report #32*, July 1975.
2. M. E. Lesk, "Lex — A Lexical Analyzer Generator", *Computing Science Technical Report #39*, October 1975.

Appendix. Suffixes and Transformation Rules

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


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


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

The following is an excerpt from the default rules file:

```
.SUFFIXES : .o .c .e .r .f .y .yr .ye .l .s
YACC=yacc
YACCR=yacc -r
YACCE=yacc -e
YFLAGS=
LEX=lex
LFLAGS=
CC=cc
AS=as -
CFLAGS=
RC=ec
RFLAGS=
EC=ec
EFLAGS=
FFLAGS=
.c.o :
    $(CC) $(CFLAGS) -c $<
.e.o .r.o .f.o :
    $(EC) $(RFLAGS) $(EFLAGS) $(FFLAGS) -c $<
.s.o :
    $(AS) -o $@ $<
.y.o :
    $(YACC) $(YFLAGS) $<
    $(CC) $(CFLAGS) -c y.tab.c
    rm y.tab.c
    mv y.tab.o $@
.y.c :
    $(YACC) $(YFLAGS) $<
    mv y.tab.c $@
```







Setting Up PWB/UNIX

R. C. Haight

Bell Laboratories
Murray Hill, New Jersey 07974

W. D. Roome

Bell Laboratories
Piscataway, New Jersey 08854

L. A. Wehr

Bell Laboratories
Murray Hill, New Jersey 07974

1. INTRODUCTION

1.1 Prerequisites

Before attempting to generate a PWB/UNIX* system, you should understand that a considerable knowledge of the attendant documentation is required and assumed. In particular, you should be very familiar with the following documents:

- *The UNIX Time-Sharing System*
- *PWB/UNIX User's Manual*
- *C Reference Manual*
- *Administrative Advice for PWB/UNIX*
- *PWB/UNIX Operations Manual*

A complete list of pertinent documentation is given in *Documents for PWB/UNIX*. Also, throughout this document, each reference of the form *name(N)*, where N is a Roman number, refers to page *name* in Section N of the *PWB/UNIX User's Manual*.

You must have a basic understanding of the operation of the hardware. This includes the console panel, the tape drives, and the disk drives, all of which are assumed to have standard addresses and interrupt vectors. If you do not understand something, consult the *PWB/UNIX Operations Manual* and DEC* maintenance manuals. It is also assumed that the hardware works and has been completely installed. The DEC diagnostics should have been run to test the configuration, and you must have a detailed description of the hardware, including device addresses, interrupt vectors, and bus levels. This information is very important to generate the PWB/UNIX system.

1.2 Procedure

PWB/UNIX is distributed on several magnetic tapes, recorded 9-track at 800bpi. Tape 1 is essential to the entire procedure, because it contains the initial load program and a copy of the *root* file system. The initial load program will copy a file system from tape (either a TU10 or a TU16) to disk (either an RP03 or an RP04). Note that RP05 and RP06 drives are considered to be equivalent to RP04 drives; any differences will be noted explicitly. Once the *root* file system has been successfully loaded to disk, PWB/UNIX may be booted and the available utility programs may then be used to complete the installation.

The remaining tapes contain source text and supplemental commands. These tapes contain essential information to generate a new system that will match your particular hardware and software environment.

In order for any of the update procedures to work correctly, you *must* be *running* PWB/UNIX. Older versions of UNIX cannot be correctly *updated* with a PWB/UNIX system. The *cpio(1)* program will not replace any file if its replacement has a modification time that is less than (i.e., earlier than) the

* UNIX is a Trademark/Service Mark of the Bell System.

modification time of the original file. This can be due to local modifications. Furthermore, certain administrative files (e.g., *passwd*, *crontab*) are sent with modification times of Jan 1, 1970 to ensure that they do not replace their counterparts during updates. Any file not copied will cause *cpio(1)* to print a message to that effect. These messages should always be investigated to ensure that any files not copied were of that type. However, note that, depending on respective modification times, a locally-modified file may get updated, thus destroying the local modifications.

There are several difficulties that can arise when installing a PWB/UNIX system. One of the most common problems is running out of disk space when performing an update. Should this occur, the original contents of the file system should be restored from a backup copy and the contents of the update tape should be read into a spare file system using the *cpio(1)* program. Unwanted material can then be removed and the original file system can be updated from this new file system using the *-p* option of *cpio(1)*. Modification times of files should also be preserved using the *-m* option of *cpio(1)*.

This document is not strictly linear. Read it thoroughly, from start to finish, and then read it again. Also, remove the write ring, if present, from every distributed tape to guard against accidental erasure.

2. LOAD PROCEDURES

2.1 Tape 1 (root) Format

Tape 1 is the initial load tape and contains four files: a loader, a physical copy of the *root* file system, the *cpio(1)* program, and a *cpio(1)* structured copy of the *root* file system. *Root* refers to the directory *"/*", which is the root of all the directory trees. The format of this tape is as follows:

record 0:	PWB tape boot loader—512-bytes;
record 1:	PWB tape boot loader—512-bytes (same as record 0);
remainder of file 1:	PWB initial load program—several 512-byte records;
end-of-file.	
file 2:	<i>root</i> file system (physical)—several 5120-byte records (blocking factor 10);
end-of-file.	
file 3:	<i>cpio</i> program—several 512-byte records;
end-of-file.	
file 4:	<i>root</i> file system (structured in <i>cpio</i> format)—several 512-byte records;
end-of-file.	
end-of-file.	

The *root (/)* file system contains the following directories:

bck:	Directory used to mount a backup file system for file restoration.
bin:	Public commands; most of what's described in Section I of the <i>PWB/UNIX User's Manual</i> .
dev:	Special files, all the devices on the system.
etc:	Administrative programs and tables.
lib:	Public libraries, parts of the assembler, C compiler.
mnt:	Directory used to mount a file system.
sys:	Directory used to mount the <i>/sys</i> file system.
tmp:	Directory used for temporary files; should be cleaned at reboot.
usr:	Directory used to mount the <i>usr</i> file system; user directories often kept here also.

2.2 Initial Load of root

Mount Tape 1 on drive 0 and position it at the load point. Next, bootstrap the tape by reading either record 0 or record 1 into memory starting at address 0 and start execution at address 0. This may be accomplished by using a standard DEC ROM bootstrap loader, a special ROM, or some manual procedure. See *romboot(VIII)*, *tapeboot(VIII)*, and *70boot(VIII)*.

The PWB/UNIX tape boot loader will then type "PWB tape boot loader" on the console terminal and read in and execute the PWB/UNIX initial load program. The program will then type detailed instructions about the operation of the program on the console terminal. Next, it will ask what type of disk drive you have and which drive you plan to use for the copy. The disk controller used must be at the standard DEC address indicated by the program. However, other disk controllers on your system may be at non-standard addresses. You must mount a formatted, error-free pack on the drive you have indicated. If necessary, use the appropriate DEC diagnostic program to format the pack. Note that the pack will be written on. Next, the program will ask what type of tape drive you have and which drive contains Tape 1. Normally, this will be drive 0, but the program will work with other drives. Note that the tape is currently positioned correctly after the end-of-file between the initial load program and the *root* file system. When everything is ready, the program will copy the file system from the tape to disk and give instructions for booting PWB/UNIX. After the copy is complete and you have booted the basic version of PWB/UNIX, *check*(VIII) the *root* file system and browse through it.

2.3 Update of root

It is very important that the system be running in *single-user* mode during the update phase. To update an already existing *root* file system, the last two files on Tape 1 will be used. It is necessary to first make a copy of your *root* file system using *volcopy*(VIII) and then update this copy. The copy should be made on a separate disk pack using the same section number as your *root* file system (almost always section 0). Also, after the update is completed, check if any of your local administrative files in the directory *etc* need modification. Most of these are mentioned in Section 4 below.

Mount Tape 1 on drive 0 and position it at the load point. We assume that disk drive 1 is available for making the copy, and that the *root* file system is on */dev/rp0*. The following procedure will then make a copy of the *root* file system, and then update this copy. Note that */dev/mt4* refers to tape drive 0 but has the side effect of spacing forward to the next end-of-file (no rewind option).

```
vcpgm root /dev/rp0 pfname1 /dev/rp10 pfname2
mount /dev/rp10 /bck
echo </dev/mt4
echo </dev/mt4
cp /dev/mt4 /bck/bin/cpio
chmod 755 /bck/bin/cpio
chown bin /bck/bin/cpio
cd /bck
/bck/bin/cpio -idm </dev/mt0
cd /
umount /dev/rp10
```

where *vcpgm* is "vc50" for RP03 disks and "vc88" for RP04/5/6, while *pfname1* and *pfname2* are the volume names of the source and destination disk packs, respectively. If the new copy is satisfactory, shut down and halt the system, change disk packs, and reboot the system using the new *root*.

2.4 Tape 2 (/usr) Format

Tape 2 contains the *usr* file system in *cpio*(I) format. The *usr* file system contains commands and files that must be available (mounted) when the system is in *multi-user* mode. The tape contains the following directories:

actg:	Accounting information.
adm:	Miscellaneous administrative command and data files, including the connect-time accounting file <i>wtmp</i> .
bin:	Public commands; an overflow for <i>lib</i> .
dict:	Dictionaries for word processing programs.
fort:	FORTRAN support files; the file <i>errors</i> is important.
games:	Various demonstration and instructional programs.
hasp:	IBM RJE programs and data.
include:	Public C language <i>#include</i> files.
jfo:	Stellar information for the <i>azel</i> (VI) program.

lib: Archive libraries, including the text processing macros; also contains data files for various programs, such as *spell*(I) and *cron*(VIII).
man: Source for the *PWB/UNIX User's Manual*; see *man*(I).
mdec: Hardware bootstrap loaders and programs.
news: Place for all the various system news.
pub: Handy public information, e.g., table of ASCII characters.
rje: More RJE information.
tmp: Directory for temporary files; should be cleaned at reboot.

A table of contents of this tape may be obtained by using the `-t` option of *cpio*(I); also after installation, files and directories deemed useless by the local administrator may be easily removed by the program *rmall*(VIII). Alternately, only parts of the tape may be extracted using the pattern matching capabilities of *cpio*(I).

2.5 Initial Load of /usr

Mount Tape 2 on drive 0 and position it at the load point. Next, a file system must be mounted as */usr*. The ultimate size and location of this file system on a device is an administrative decision; initially, the following procedure will suffice:

```
cd /
mkfs /dev/rpl 35000
labelit /dev/rpl usr pkname
icheck -sflag /dev/rrpl
check /dev/rrpl
mount /dev/rpl /usr
cd /usr
cpio -idm </dev/mt0
```

where *pkname* is the volume name of the pack (e.g., "p0001"), while *flag* is "3" for RP03 disks and "4" for RP04/5/6.

Since */usr* must be mounted when the system is in *multi-user* mode, the file */etc/rc* must be changed to include the command line:

```
mount /dev/rpl /usr
```

This line should be appended after the first line of */etc/rc*, which initializes the mount table *mnttab*(V). Next the file */etc/checklist* should be changed to include */dev/rrpl*. See also *check*(VIII), *icheck*(VIII), *labelit*(VIII), *mkfs*(VIII), *mount*(VIII).

2.6 Update of /usr

It is advisable that the system be running in *single-user* mode during the update phase. It is also wise to first make a copy of your */usr* file system for backup purposes. Next, mount Tape 2 on drive 0 and position it at the load point. The */usr* file system *must* also be mounted. The following procedure will perform the update:

```
cd /usr
cpio -idm </dev/mt0
```

2.7 Tape 3 (/sys) Format

Tape 3 contains the */sys* file system in *cpio*(I) format. The */sys* file system contains the source program text for the operating system and for the public commands. File names that end in *.c* are in C, *.s* are in assembler, and those that end in *.a* are archive files. The tape contains the following directories:

c: Source for the C language compiler.
fort: Source for the FORTRAN language compiler.
source: Source for the public commands and libraries, including sub-directories:
lex: LEX program
plot: Plotting routines

rje: IBM RJE programs
s1: Commands [a-l]*
s2: Commands [m-z]*
s3: Assembler library
s4: C library
s5: PWB library
s7: Word processing programs
s8: Administrative programs
sccs4: Source Code Control System commands
sno: SNOBOL compiler.

sys: Source for the operating system, including sub-directories:

cf: Configuration information
hd: Header files
io: I/O handlers
ml: Machine-level routines
os: Operating system.

A table of contents of this tape may be obtained by using the `-t` option of `cpio(1)`; after installation, files and directories deemed useless by the local administrator may be easily removed by the program `rmall(VIII)`. Alternately, only parts of the tape may be extracted using the pattern matching capabilities of `cpio(1)`. To generate a new PWB/UNIX operating system, only the `sys` directory is needed.

2.8 Initial Load of /sys

Mount Tape 3 on drive 0 and position it at the load point. Next, a file system must be mounted as `/sys`. The ultimate size and location of this file system is an administrative decision; initially, the following procedure will suffice:

```
cd /
mkfs /dev/rp2 35000
labelit /dev/rp2 sys pkname
icheck -sflag /dev/rp2
check /dev/rp2
mount /dev/rp2 /sys
cd /sys
cpio -idm </dev/mt0
```

where *pkname* is the volume name of the pack (e.g., "p0001"), while *flag* is "3" for RP03 disks and "4" for RP04/5/6.

The `/sys` file system need not be mounted when the system is in *multi-user* mode; however, it must be mounted in order to reconfigure your system. If you want to always mount the `/sys` file system, include the following line in `/etc/rc` after the line which mounts the `/usr` file system:

```
mount /dev/rp2 /sys
```

After making a new system, you may want to `umount(VIII)` the `/sys` file system and save it on a new tape with `volcopy(VIII)`.

2.9 Update of /sys

It is advisable that the system be running in *single-user* mode during the update phase. It is also wise to first make a copy of your `/sys` file system for backup purposes. Next, mount Tape 3 on drive 0 and position it at the load point. The `/sys` file system *must* also be mounted. The following procedure will perform the update:

```
cd /sys
cpio -idm </dev/mt0
```

2.10 Tape 4 (troff and eqn) Format

Tape 4 (if supplied) contains the *troff* and *eqn* software in *cpio*(1) format. The tape contains the following directories:

- croff: Source for the *troff*(1) program including the sub-directory:
 - font: font tables.
- eqn: Source for the *eqn*(1) program.

A table of contents of this tape may be obtained by using the *-t* option of *cpio*(1).

2.11 Initial Load or Update of troff and eqn Programs

It is advisable that the system be running in *single-user* mode during the update phase. Mount Tape 4 on drive 0 and position it at the load point. Make sure that the */sys* and */usr* file systems are mounted. The *make*(1) command will compile and install all of the necessary material. The following procedure should be used:

```
cd /sys/source/s7
cpio -idm </dev/mt0
cd /usr/bin
make -f .makefile4
cd /usr/lib
make -f .makefile2
```

3. CONFIGURATION PLANNING

3.1 PWB/UNIX Configuration

The basic PWB/UNIX operating system supplied on Tape 1 supports only the console, a disk controller (disk drive 0), and a tape controller (tape drive 0). The actual configuration of your system must be described by you. All of the PWB/UNIX operating system source code and object libraries are in */sys/sys*. All of the configuration information is kept in the directory */sys/sys/cf*. There are only two files that must be changed to reflect your system configuration, *low.s* and *conf.c*; the program *config*(VIII) makes this task relatively simple.

Config requires a *system description file* and produces the two needed files. The first part of the system description file lists all of the hardware devices on your system. Next, various system information is listed. A brief explanation of this information follows. For more details of syntax and structure, see *config*(VIII).

- *root*—Specifies the device where the *root* file system is to be found. The device must be a block device with read/write capability since this device will be mounted read/write as *"/*". Thus, a tape can not be mounted as the *root*, but can be mounted as some read-only file system. Normally, *root* is disk drive 0, section 0.
- *swap*—Specifies the device and blocks that will be used for *swapping*. *Swplo* is the first block number used and *nswap* indicates how many blocks, starting at *swplo*, to use. *Swplo* can not be zero. Typically, systems require approximately 2,000 blocks. Care must be taken that the swap area specified does not overlap any file system. For example, if section 0 is 8,000 blocks long, the *root* file system could occupy the first 6,000 blocks and *swap* the remaining 2,000 by specifying:

```
root rp04 0
swap rp04 0 6000 2000
```

- *dump*—Specifies the device to be used to dump memory after a system crash. Currently only the TU10 and TU16 tape drives are supported for this purpose.
- *buffers*—Specifies how many *system buffers* to allocate. One buffer is needed for every mounted file system, and various I/O handlers temporarily claim several buffers. In general, you will want as many buffers as possible without exceeding the system space limitations. Normally, *butlers* is in the range of 30-50. Each entry requires 538 bytes.

- *inodes*—Specifies how many *inode table* entries to allocate. Each entry represents a unique open inode. When the table overflows, the warning message “Inode table overflow” will be printed on the console. The table size should be increased if this happens regularly. The number of entries used depends on the number of active processes, texts, and mounts. Normally, *inodes* is in the range of 150-300. Each entry requires 32 bytes.
- *files*—Specifies how many *open-file table* entries to allocate. Each entry represents an open file. When the table overflows, the warning message “no file” will be printed on the console. The table size should be increased if this happens regularly. Normally, *files* is in the same range as the number of inodes. Each entry requires 8 bytes.
- *mounts*—Specifies how many *mount table* entries to allocate. Each entry represents a mounted file system. The *root (/)* will always be the first entry. When full, the *mount(II)* syscall will return the error EBUSY. Normally, *mounts* is in the range of 8-16. Each entry requires 6 bytes.
- *coremap*—Specifies how many entries to allocate to the *list of free memory*. Each entry represents a contiguous group of 64-byte blocks of free memory. When the list overflows, due to excessive fragmentation, the system will undoubtedly crash in an unpredictable manner. The number of entries used depends on the number of processes active, their sizes, and the amount of memory available. Normally, *coremap* is in the range of 50-100. Each entry requires 4 bytes.
- *swapmap*—Specifies how many entries to allocate to the *list of free swap blocks*. Exactly like the *coremap*, except it represents free blocks in the swap area, in 512-byte units. Each entry requires 4 bytes.
- *calls*—Specifies how many *callout table* entries to allocate. Each entry represents a function to be invoked at a later time by the clock handler. The time unit is 1/60 of a second. The callout table is used by the terminal handlers to provide terminal delays and by various other I/O handlers. When the table overflows, the system will crash and print the panic message “Timeout table overflow” on the console. Normally, *calls* is in the range of 30-60. Each entry requires 6 bytes.
- *procs*—Specifies how many *process table* entries to allocate. Each entry represents an active process. The scheduler is always the first entry and *mt(VIII)* is always the second entry. The number of entries depends on the number of terminal lines available and the number of processes spawned by each user. The average number of processes per user is in the range of 2-5. When full, the *fork(II)* syscall will return the error EAGAIN. Normally, *procs* is in the range of 50-200. Each entry requires 28 bytes.
- *texts*—Specifies how many *text table* entries to allocate. Each entry represents an active read-only text segment. Such programs are created by using the *-i* or *-n* option of the loader *ld(I)*. When the table overflows, the warning message “out of text” is printed on the console. Normally, *texts* is in the range of 25-50. Each entry requires 12 bytes.
- *chists*—Specifies how many *character list buffers* to allocate. Each buffer contains up to six bytes. The buffers are dynamically linked together to form input and output queues for the terminal lines and various other slow-speed devices. The average number of buffers needed per terminal line is in the range of 5-10. When full, input characters from terminals will be lost and not echoed. Normally, *chists* is in the range of 100-300. Each entry requires 8 bytes.

3.2 PWB/UNIX Generation

To generate a new PWB/UNIX operating system, mount the */sys* file system and follow this general procedure:

```
cd /sys/sys/cf
ed dfile
a
  {information as described above}
.
w
q
config dfile
make NAME=name
```

Dfile is the name of the file containing the configuration information, and *name* is a character string of at most eight characters used to identify a specific system (such as "pwba0501"). The procedure will assemble *low.s*, compile *conf.c*, and load them together with the object libraries into a file called *name*.

The system has a finite address space, so that if table sizes or the number of device types are too large, various error messages will result and the above procedure will only create an *a.out* file. In particular, the maximum available data space is 49,152 bytes. The actual data space requested can be found by using *size(1)* on *a.out* and adding the *data* and *bss* segment sizes. One then reduces the specified values for the various system entries (normally, the number of buffers) until it all fits. The amount of space in the *bss* segment used for each entry is indicated in Section 3.1 above.

When you are satisfied with the new system, you can test it by the following procedure:

```
cd /sys/sys/cf
cp name /
cd /
rm /unix
ln /name /unix
sync
```

Then halt the processor and reboot *unix*. Note that this procedure results in two names for the operating system object, the generic *linux*, and the actual name, say *lpwba0501*. An old system may be booted by referring to the actual name, but remember that many programs use the generic name *linux* to obtain the *name-list* of the system.

If the new system does not work, verify that the correct device addresses and interrupt vectors have been specified. If the *wrong* interrupt vector and the *correct* device address have been specified for a device, the operating system will print the error message "stray interrupt at XXX" when the device is accessed, where XXX is the correct interrupt vector. If the *wrong* device address is specified, the system will crash with a panic trap of type 0 (indicating a UNIBUS® timeout) when the device is accessed.

3.3 Special Files

A special file must be made for every device on your system. Normally, all special files are located in the directory */dev*. Initially, this directory will contain:

<i>tty8</i>	console terminal
<i>mem</i> , <i>kmem</i> , <i>null</i>	see <i>mem(IV)</i>
<i>tty</i>	see <i>tty(IV)</i>
<i>rp[0-7]</i> , <i>rrp[0-7]</i>	disk drive 0, sections 0-7
<i>mt0</i> , <i>rmt0</i>	tape drive 0 (800 bpi)
<i>mt4</i> , <i>rmt4</i>	tape drive 0 (800 bpi, no rewind).

There are two types of special files, block and character. This is indicated by the character *b* or *c* in the listing produced by *ls(1)* with the "-l" flag.

In addition, each special file has a major device number and a minor device number. The major device number refers to the device type and is used as an index into either the *bdevsw* or *cdevsw* table in the configuration file *conf.c*. The minor device number refers to a particular unit of the device type and is used only by the driver for that type. For example, using the following sample portion of a configuration file, a block special file with major device number 1 and minor device number 0 would refer to the TU10 magtape, drive 0, while a character special file with major device number 1 and minor device number 4 would refer to the DH11 asynchronous multiplexor, line 4.

```
int      (*bdevsw[])(
{
/* 0*/  &hpopen, &hpclose, &hpstrategy, &hptab,
/* 1*/  &tmopen, &tmclose, &tmstrategy, &tmtab,
};
int      (*cdevsw[])(
{
/* 0*/  &klopen, &kfclose, &khread, &kfwrite, &kls/tty,
/* 1*/  &dhopen, &dclose, &dhread, &dhwwrite, &dhs/tty,
/* 2*/  &nulldev, &nulldev, &mmread, &mmwrite, &nodev,
/* 3*/  &nodev, &nodev, &nodev, &nodev, &nodev,
/* 4*/  &nodev, &nodev, &nodev, &nodev, &nodev,
/* 5*/  &nodev, &nodev, &nodev, &nodev, &nodev,
/* 6*/  &tmopen, &tmclose, &tmread, &tmwrite, &nodev,
/* 7*/  &hpopen, &hpclose, &hpread, &hpwrite, &nodev,
/* 8*/  &nodev, &nodev, &nodev, &nodev, &nodev,
/* 9*/  &nodev, &nodev, &nodev, &nodev, &nodev,
/*10*/  &nodev, &nodev, &nodev, &nodev, &nodev,
/*11*/  &nodev, &nodev, &nodev, &nodev, &nodev,
/*12*/  &nodev, &nodev, &nodev, &nodev, &nodev,
/*13*/  &syopen, &nulldev, &syread, &sywrite, &sys/tty,
};
```

The program *mknod*(VIII) creates special files. For example, the following would create *part* of the initially-supplied */dev* directory:

```
cd /dev
mknod tty8 c 0 0
mknod mem c 2 0; mknod kmem c 2 1; mknod null c 2 2
mknod tty c 13 0
mknod rp0 b 0 0; mknod rrp0 c 7 0
mknod mt0 b 1 0; mknod rmt0 c 6 0
mknod mt4 b 1 4; mknod rmt4 c 6 4
```

After the special files have been made, their access modes should be changed to appropriate values by *chmod*(1). For example:

```
cd /dev
chmod 622 tty8
chmod 644 mem kmem
chmod 666 null
chmod 666 tty
chmod 400 rp0 rrp0
chmod 666 mt0 rmt0
chmod 666 mt4 rmt4
```

Note that file names have no meaning to the *operating system* itself; only the major and minor device numbers are important. However, many *programs* expect that a particular file is a certain device. Thus, by convention, special files are named as follows:

<i>block device</i>	<i>conf.c</i>	<i>/dev</i>
RP03 disk	rp	rp*
RP04/5/6 disk	hp	rp*
RS03/4 fixed head disk	hs	rs*
TU10 tape	tm	mt*
TU16 tape	ht	mt*

<i>character device</i>	<i>conf:c</i>	<i>ldev</i>
DL11 asynch line	kl	tty?
DH11 asynch line mux	dh	tty?
DN11 auto call unit	dn	dn*
DU11 synch line	du	du*
DQS11B synch line	dqs	rjei
LP11 line printer	lp	lp*
RP03 disk	rp	rrp*
RP04/5/6 disk	hp	rrp*
RS03/4 fixed head disk	hs	rrs*
TU10 tape	tm	rmt*
TU16 tape	ht	rmt*
memory	mm	mem, kmem, null
terminal	sy	tty

For those devices with a *ldev* name ending in "*", this character is replaced by a string of digits representing the *minor* device number. For example:

```
mknod /dev/mt1 b 1 1
mknod /dev/rp24 b 0 024
```

For those devices with a *ldev* name ending in "?", this character is replaced by any alphanumeric character, except "x". For example:

```
mknod /dev/ttya c 1 0
mknod /dev/tty8 c 0 0
```

Note that *ty8* is historically the console, and that for disks, an octal number scheme is maintained because each drive is split eight ways. Thus, *ldev/rp24* refers to section 4 of physical drive 2. There is also a special file, *ldev/swap*, that is used by the program *ps(1)*. This file must reflect what device is used for swapping and must be readable. For example:

```
rm /dev/swap
mknod /dev/swap b 0 0
chmod 444 /dev/swap
```

3.4 File Systems

The maximum number of blocks in a file system is 65,535. However, the RP03/4/5/6 disks contain many more blocks. Therefore, each physical pack is split into eight logical sections. This split is defined in the operating system by a table with eight entries. Each table entry is two words long. The first specifies how many blocks are in the section, the second specifies the starting cylinder. Initially, these are set as follows:

<i>section</i>	<i>RP03</i>		<i>RP04/5</i>		<i>RP06</i>	
	<i>count</i>	<i>start</i>	<i>count</i>	<i>start</i>	<i>count</i>	<i>start</i>
0	7600	0	11286	0	11286	0
1	36200	38	53504	27	65535	27
2	36200	219	53504	155	53504	155
3	65535	40	53504	283	53504	283
4	36200	22	65535	27	65535	184
5	40600	203	65535	184	65535	341
6	0	0	29260	341	65535	498
7	0	0	0	0	65535	655

These values are described to the system in the header file */sys/sys/hdllo.h* which may be changed by using the editor *ed(1)*. After such a change, the system must be made again (see Section 3.2).

A file system starts at block 0 of a section of the disk and may be as large as the size of that section; if it is smaller than the size of a section, the remainder of that section is unused. Note that the sections themselves may overlap physical areas of the pack, but the file systems must *never* overlap.

The program *mkfs*(VIII) is used to initialize a section of the disk to be a file system. Next, the program *labelit*(VIII) is used to label the file system with its name and the name of the pack. Finally, the file system may be checked for consistency by using *check*(VIII) and the free list reorganized by using *ichk*(VIII). The file system may then be mounted using *mount*(VIII).

4. ADMINISTRATIVE FILES

4.1 /etc/motd

This file contains the *message-of-the-day*. It is printed by *login*(1) after every successful *login*.

4.2 /etc/rc

On the transition from *single-user* to *multi-user* mode, *letclinit* invokes */bin/sh* with input from *letclrc* before forking to create the *multi-user* environment. That file usually contains commands to remove temporary files, mount file systems, and start daemons. For example, a typical *letclrc* would be:

```
rm /tmp/*
devnm / | /etc/setmnt
mount /dev/rpl /usr
```

This file must be edited to suit local conditions; see also *init*(VIII).

4.3 /etc/tty8

This file indicates to *letclinit* how many processes to create for the *multi-user* environment. Initially, this file will contain only:

```
18-
```

This indicates that the console terminal, namely */dev/tty8*, should have a *login* process attached when in *multi-user* mode. If this line were not present, you could not *login*(1) from the console when in that mode.

For example, to attach a *login* process to line */dev/ttya* for use by 300/150/110 baud terminals, add the following line:

```
1a0
```

Again, this file must be edited for local conditions. See *init*(VIII), *tty*(V).

4.4 /etc/passwd

This file is used to describe each *user* to the system. You must add a new line for each new user. Each line has six fields separated by colons:

1. Login name:
Normally 1-6 characters, first character alphabetic, rest alphanumeric, no upper-case.
2. Encrypted password:
Initially null, filled in by *passwd*(1). The encrypted password contains 11 bytes, while the actual password is limited to a maximum of 8 bytes.
3. User ID:
A number between 0 and 255; 0 indicates the super-user. Four other IDs are reserved:
bin::2: software administration;
adm::3: system administration;
rje::68: remote job entry administration;
games::180: miscellaneous; never a real user.
4. Accounting information:
This field is used by various accounting programs. It usually contains the user name, department number, and account number.

5. Login directory:

Full pathname that must not exceed 21 characters.

6. Program name:

If null, */bin/sh* is invoked after successful *login*. If present, the named program is invoked in place of */bin/sh*.

For example:

```
ghh::38:6824-G.H.Hurtz(4357):/usr/ghh:  
grk::26:6510-S.P.LeName(4466):/usr/grk:/bin/rsh
```

See also *passwd(V)*, *login(I)*, *passwd(I)*.

4.5 */etc/group*

This file is used to describe each *group* to the system. Each line has four fields separated by colons:

```
group name;  
encrypted password;  
numerical group ID;  
list of all login names in the group, separated by commas.
```

See also *group(V)*.

4.6 */etc/checklist*

This file contains a list of default devices to be checked for consistency by the *check(VIII)* program. The devices normally correspond to those mounted when the system is in *multi-user* mode. For example, in conjunction with the sample *etc/rc* of Section 4.2 above, and with the *root* file system mounted on */dev/rp0*, *etc/checklist* would be:

```
/dev/rrp0  
/dev/rrp1
```

4.7 */etc/startup*

This file is a secondary initialization file. It normally contains commands to initialize the various RJE systems and start the *cron(VIII)* daemon.

4.8 */etc/shutdown*

This file contains procedures to gracefully shut down the system in preparation for filesave or scheduled downtime.

4.9 */etc/filesave.?*

This file contains the detailed procedures for the local filesave.

4.10 */etc/sha*

This file contains the shell accounting information. For detailed information, see *sha(V)*.

5. REGENERATING SYSTEM DIRECTORIES

5.1 General Procedure

There are several system directories that contain various commands and archive libraries. These include */bin*, */usr/bin*, */lib*, */usr/lib*, *etc*, */usr/hasp*, and */usr/fort*. To facilitate the regeneration of these directories, the *make(I)* command is used. Each directory contains the *make* description files used to build, update, or regenerate the directory. If one description file suffices, it is named *.makefile*; otherwise, the files are named *.makefile1*, *.makefile2*, ...

For this entire procedure to work, you must be logged in as "root" (or as super-user), and the */sys* and */usr* file systems must be mounted.

For example, the following will recompile any */bin* commands that are out-of-date with respect to their corresponding source files:

```
cd /bin
make -f .makefile1
make -f .makefile2
make -f .makefile3
```

As another example, the following forces the recompilation of any */bin* commands that call *ctime*(III), in addition to recompiling anything that is out-of-date in */bin*:

```
cd /bin
make -f .makefile1 CTIME=RC
make -f .makefile2 CTIME=RC
make -f .makefile3 CTIME=RC
```

Section 5.2 on “Time Zones” explains why this is useful. Similar “...=RC” arguments force recompilation of anything that uses the C compiler (CCDEP=RC), the assembler (ASDEP=RC), the *yacc*(I) processor (YACCDEP=RC), the “-lS” library (LSDEP=RC), the “-lPW” library (LPWDEP=RC), the “-lpw” library (LpwDEP=RC), and the PWB/UNIX operating system “include” files (SYSDEP=RC).

5.2 Time Zones

Commands such as *ls*(I) and *date*(I) give Eastern times (EST or EDT). This has been known to annoy non-East Coast (USA) and non-USA users. To get your local time, you must change *ctime*(III), which converts the GMT *time*(II) time to local time, and then regenerate all commands that use *ctime*. To be precise, you must:

- Change *ctime*(III). The source currently lives in */sys/source/s4/ctime.c*, and the object resides in the standard C library, */lib/libc.a*.
- Regenerate the standard C library, as follows:

```
cd /lib
make -f .makefile1
make -f .makefile2
make -f .makefile3
make -f .makefile4
```

This works because, as a result of the change to the source for *ctime*(III), */lib/libc.a* is now out-of-date with respect to its source, namely the *ctime* part.

- Regenerate all directories, using the “CTIME=RC” option on the *make*(I) command line:

```
cd /bin
make -f .makefile1 CTIME=RC
make -f .makefile2 CTIME=RC
make -f .makefile3 CTIME=RC
cd /usr/bin
make -f .makefile1 CTIME=RC
make -f .makefile2 CTIME=RC
make -f .makefile3 CTIME=RC
make -f .makefile4 CTIME=RC
cd /lib
make -f .makefile1 CTIME=RC
make -f .makefile2 CTIME=RC
make -f .makefile3 CTIME=RC
make -f .makefile4 CTIME=RC
cd /usr/lib
make -f .makefile1 CTIME=RC
make -f .makefile2 CTIME=RC
cd /etc
make -f .makefile CTIME=RC
cd /usr/hasp
make -f .makefile CTIME=RC
cd /usr/fort
make -f .makefile CTIME=RC
```

The *current* full directory list is given above for completeness. Some directories might not have anything that uses *ctime*. Of course, you should test your new *ctime* before installing it in the standard C library.

Unfortunately, *nroff*(I) and *troff*(I) are fixed to Eastern Standard Time; they use their own conversion routines, instead of calling *ctime*(III). Fortunately, they use just the date, not the time-of-day.

5.3 /bin

Regeneration of */bin* is straightforward, except for *as*(I), *ld*(I), and *cc*(I). The *make* description files for */bin* do *not* automatically regenerate *cc*; instead, a separate procedure is used (see “*cc*” below).

As and *ld* are necessary to regenerate themselves. That is, one cannot re-assemble *as* unless a working *as* already exists. Thus, if something goes wrong and a bad *as* or *ld* is generated, you’ll have to retrieve a working version from a backup tape.

The moral: before regenerating */bin*, be sure that you either have a good backup tape, or else have emergency copies of *as* and *ld* somewhere.

5.4 /usr/bin

.makefile1 and .makefile2 regenerate the basic *usr/bin* commands; .makefile3 regenerates the SCCS commands; .makefile4 regenerates *eqn*(I) and *troff*(I). If you do not have the PWB/UNIX *troff*(I) and *eqn*(I) package (Tape 4), .makefile4 will fail.

5.5 /lib

As with *as* and *ld* in */bin*, *as2* in */lib* (“pass 2” of *as*) can’t be regenerated if it doesn’t already exist. Also, *c0*, *c1*, *c2*, and *cpp* in */lib* are parts of the C compiler, and the *make* description files for */lib* do *not* automatically regenerate them (see “*cc*” below).

5.6 /usr/lib

.makefile1 regenerates the basic *usr/lib* modules; .makefile2 regenerates tables used by *eqn*(I) and *troff*(I). If you do not have the PWB/UNIX *troff*(I) and *eqn*(I) package (Tape 4), .makefile2 will fail.

5.7 cc

/bin/cc is just the tip of the C compiler iceberg. The real work is done by *c0*, *c1*, *c2*, and *cpp* in *//ib*. Since these programs work together, they must be updated carefully. Furthermore, since the C compiler is written in C, if you install a bad C compiler you'll have to retrieve a working version from a backup tape.

Therefore, the *make* procedures for */bin* and *//ib* do not attempt to regenerate the C compiler. Instead, you must do it explicitly, by:

```
cd /sys/c/c
make -f makefile install clean
```

That creates *cc*, *c0*, *c1*, *c2*, and *cpp* in */sys/c/c*, and, if no errors occur, moves them to */bin* and *//ib*. "Clean" just removes any object files (*.o*) left around. Note that the *make* description file for the C compiler is *makefile* (no leading *.*). Only the *make* description files that live in object directories (such as */bin*) start with *.*.

Often it is better to install a new version of the C compiler as *ncc*, so that it can be tested without destroying the old (working) version. The following creates and installs a complete new version of the C compiler as *ncc*:

```
cd /sys/c/c
make -f makefile PREF=n install clean
```

As before, this creates *cc*, *c0*, *c1*, *c2*, and *cpp* in */sys/c/c*. But they are installed as *ncc*, *nc0*, *nc1*, *nc2*, and *ncpp* in */bin* and *//ib*. This works because the *cc* driver looks at the name under which it was invoked. If the first character of the name is 'c', it calls *c0*, etc., in *//ib*. But if the first character isn't 'c', it prepends that character to the names of the load modules that it calls. Thus *ncc* calls *nc0*, *nc1*, *nc2*, and *ncpp* in *//ib*.

5.8 /usr/hasp

The *.makefile* for */usr/hasp* specifies whether your RJE system supports one or two lines. The following generates a two-line system:

```
cd /usr/hasp
make -f .makefile NHASP=2
```

while:

```
cd /usr/hasp
make -f .makefile NHASP=
```

generates a one-line system (the default is one-line). To change the default, change the "NHASP=" line near the beginning of the *.makefile*.

5.9 Changes to Operating System Parameters

A few commands, such as *ps(I)* and *adb(I)*, depend on some of the parameters embedded in the PWB/UNIX operating system, such as the maximum number of processes and the size of the swap area (see Section 3.1). If these parameters are changed, you should recompile all the commands that depend on the PWB/UNIX operating system "include" files. To do so, first generate the new PWB/UNIX operating system, and then execute the following:

```
cd /etc
make -f .makefile SYSDEP=RC
cd /bin
make -f .makefile1 SYSDEP=RC
make -f .makefile2 SYSDEP=RC
make -f .makefile3 SYSDEP=RC
cd /usr/bin
make -f .makefile1 SYSDEP=RC
make -f .makefile2 SYSDEP=RC
make -f .makefile3 SYSDEP=RC
make -f .makefile4 SYSDEP=RC
```

The “SYSDEP=RC” option on the *make*(1) command line forces recompilation of anything that depends on the operating system “include” files. The other system directories do not have any commands that depend on the operating system parameters.

The above is only necessary when PWB/UNIX operating system *parameters* are changed. It is *not* necessary if you have only added a device or changed device addresses.


5.10 Miscellaneous


You should be familiar with *make*(1), *regen*(VIII), and *Make—A Program for Maintaining Computer Programs* by S. I. Feldman.

The *make* diagnostic “Don’t know how to make xxx” means that the source file xxx doesn’t exist.

6. ACKNOWLEDGEMENTS

We thank M. J. Petrella and M. F. Pucci for verifying, testing, and improving the installation procedures described in this document. We also thank M. J. Petrella for his help in preparing this memorandum, for updating it to reflect the latest available information, and for editing it into its present form.





Administrative Advice for PWB/UNIX

R. C. Haight

Bell Laboratories
Murray Hill, New Jersey 07974

The material presented here is based on the author's experiences and opinions. Nevertheless, it may prove useful. The material on phototypesetting was contributed by D. W. Smith.

1. SYSTEM CAPACITY

The figures below are approximations based on our experience over several years.

<i>Hardware Configuration</i>	<i>Number of Simultaneous Users</i>
PDP-11/45; 124K-word memory; RP03 disk†	16
Above with RP04 disk†	20
Above with cache memory buffer	25
PDP-11/70; 128K-word memory; RP04 (RP05, RP06) disks† (2 or more drives)	32
Above with 256K-word memory and a disk drive (or fixed-head disk) set aside for the root file system	48

† Or equivalent.

2. DISK FREE SPACE

Making files is easy under PWB/UNIX.* In addition, the PWB/UNIX RJE system with a 50K-baud link can theoretically consume 40,000 blocks/hour. If the free inode count falls below 100, the system spends most of its time rebuilding the free inode array. If disk free space falls below 1,200 blocks in the */usr* file system, RJE shuts down completely (it starts complaining and stops receiving output earlier than that). If a file system runs out of space (and the console switches are non-zero), the system prints "no-space" messages and does little else. To avoid problems, the following start-of-day free counts should be maintained:

- The file system containing *"/tmp"* (temporary files):
 - 16-user system: 1,500 free blocks.
 - 48-user system: 3,000 free blocks.
- The file system containing *"/usr"* (RJE and accounting files):
 - 3,000 to 6,000 free blocks, depending on load.
- User file systems:
 - 6% to 10% free, depending on user habits (2,000 blocks minimum).

3. GET A SPARE DISK DRIVE

- Without a spare disk drive, the system will be *down* when a drive is *down*.
- Without a spare drive, it is difficult to reorganize file systems and to restore user files.

4. DO NOT FOOL AROUND WITH IMPERFECT DISK PACKS

- Buy only 100% certified packs and test them.
- If a pack develops errors, recondition it, or get rid of it.

* UNIX is a Trademark of Bell Laboratories.

5. CREATING FILE SYSTEMS

The *mkfs*(VIII) program creates new file systems. The *icheck*(VIII) program optimizes the free list organization and posts the free counts in the super-block. *Labelit* (see *volcopy*(VIII)) records the file system name and volume-id in the super-block. The following command sequence readies a new file system (on a new RP04 or RP05 pack) for mounting:

```
mkfs /dev/rpX 65000 (where "X" is the minor device number)
labelit /dev/rpX name pack#
icheck -s4 /dev/rpX
```

See *Setting up PWB/UNIX* by R. C. Haight, W. D. Roome, and L. A. Wehr for more information.

6. PROTECTING USER FILES

Users, especially inexperienced ones, occasionally remove their own files. Open files are sometimes lost when the system crashes. Once in a great while, an entire file system will be destroyed (picture a disk controller that goes bad and writes when it should read). Here is a suggested file backup procedure:

- Each day, copy all user file systems to backup packs. Keep these packs 3 to 5 days before re-using them.
- Once a week, copy each file system to tape. Keep weekly tapes for 8 weeks.
- Keep bi-monthly tapes "forever" (they should be re-copied at least once a year).

The most recent weekly tapes should be kept off-premises. The other tapes should be in a fire-proof safe.

When PWB/UNIX goes down, active files can get scrambled. Your users will not want to start the day over every time your system fails. In addition to good backup, you *must* have file-system patching expertise available (on-site or on-call). If you ever re-boot the system for general use without checking out the file systems, terrible things will happen (we once had five duplicate entries on a file-system free list—this ruined over 100 new files in just three days). See *Repairing Damaged PWB/UNIX File Systems* by P. D. Wandzilak and the "Boot Procedures" Section of the *PWB/UNIX Operations Manual* by M. E. Pearlman.

7. CONTROLLING DISK USAGE

If your PWB/UNIX system is a success, you will soon run out of disk space:

- During the considerable delay before you can get more drives, you will need to control usage:
 - The *du*(1) command should be executed (after hours) regularly (e.g., weekly) and the output kept (in an accessible file) for later comparison. In this way you can spot users who are rapidly increasing their disk usage.
 - The *find*(1) can be used to locate inactive (or large) files. Example:

```
find / -mtime +90 -a -atime +90 -a -print >somefile
```

writes in "somefile" the names of files neither written nor accessed in the last 90 days. Of course, this works best if you are super-user.

- You will also have to balance usage between file systems. To do this you will have to move user directories. Users should be taught to accept file system name changes (and to program around them—preferably ahead of time). The user's login directory name (available in the Shell variable "\$S" or from the subroutine "logdir()") should be utilized to minimize pathname dependencies. User groups with more extensive file system structures should use the Shell variable "\$w" (file system name).
- The *find*(1) and *cpio*(1) command can be used to move user directories and to manipulate the file system tree. The following sequence is useful (it moves, via magnetic tape, the directory tree *userX* from file system *filesys1* to file system *filesys2*):


```
cd /filesys1/userx
* du -a >/tmp/x
find . -print | cpio -o >/dev/mt0
mkdir /filesys2/userx
chown userx /filesys2/userx
chdir /filesys2/userx
cpio -idm </dev/mt0
* du -a >/tmp/y
* diff /tmp/x /tmp/y
    (only directory sizes should differ)
* rm /tmp/x /tmp/y
rmail -r -d -f /filesys1/userx
```

Starred items are (optional) precautions. When moving more than one user in this way:

- Keep users with common interests in the same file system (they may have linked files).
- Move groups of users who may have linked files with a single *cpio* (otherwise linked files will be unlinked and duplicated).

8. REORGANIZING FILE SYSTEMS

The procedures outlined in sections “Creating File Systems” and “Controlling Disk Usage” above can be combined to reorganize whole file systems. Reorganization can improve system response time. This is particularly true of the *root* file system (which must be reorganized with all other file systems unmounted) and *usr*.

9. KEEP DIRECTORY FILES SMALL

Directories larger than 4K bytes (256 entries) are very inefficient because of file system indirection. A PWB/UNIX user once complained that it took the system ten minutes to complete the login process. It turned out that his login directory was 25K bytes long, and the login program spent that time fruitlessly looking for a non-existent “.mail” file.

10. ADMINISTRATIVE USE OF “CRON”

The program *cron*(VIII) is useful in the administration of the system; it can be used to:

- Turn off the programs in directory *usr/games* during prime time.
- Recognize holidays (for accounting purposes).
- Run programs off-hours:
 - accounting;
 - file system administration;
 - long-running, user-written Shell procedures (using the *setuid*(VIII) command), for example:

```
setuid userx /filesysA/userx % userx_shell arg ...
```

11. ALLOCATING RESOURCES TO USERS

A prospective user should obtain connect-time and file-space authorization through appropriate channels. Once this is done, the user should apply for a login by providing the following information to the “system administrator”:

- User’s name.
- Suggested login name (not more than 8 characters, beginning with a lower-case letter).
- Initial password (yes, everyone should have one—not more than 8 characters, but more than 5, and *not* in Webster’s Unabridged).
- Relationships to other users and RJE requirements (this influences the choice of the file system).
- Estimate of required file space (this also influences the choice of the file system).

If you plan to support a large user community, you should ration user-id numbers (256 is a small number). Several login names per user-id will work for users working on the same project.

12. THE MATTER OF ACCOUNTING AND USAGE

You should run the accounting programs even if you do not "bill" for service. Otherwise, your users' habits (especially *bad* habits) will be a mystery to you. For instance, you can discover unused login names by extracting, from the daily accounting data, the date of last login for each login name.

13. DIAL LINE UTILIZATION

If prime-time dial line utilization gets much over 70%, users will start to encounter busy signals when dialing in. This, in turn, will lead to "line hogging." The only solutions are to get a larger (another) machine, or to get rid of users. Manual policing will help some, but "automatic" policing will *invariably* be subverted by users.

14. "BIRD-DOGGING"

When the system is busy (lines busy and/or slow response), someone should determine why this is so. The *who(1)* command lists the people logged in. The *ps(1)* command shows what they are doing. (The undocumented *letc/whodo* command combines the output of *who* and *ps*.) Unfortunately, *ps* operates from heuristics that can consistently fail to report certain processes in a busy system. That is, one must be careful about hanging up an apparently inactive line. The *lastcom(VIII)* command reads the Shell accounting file *letc/sha* backwards from the most recent entry. It will print entries for selected lines (tty identification characters) or login names.

15. 300-BAUD TERMINALS

Upper-case-only terminals are a pain on PWB/UNIX. Get full-duplex, full-ASCII terminals. Hardware horizontal tabbing is very desirable, because it increases output speed and lowers system overhead. A fair proportion of your terminals should provide for correspondence-quality hard-copy output to take advantage of the PWB/UNIX word-processing capabilities; see *terminals(VII)*.

16. BE CAREFUL OF HIGH-SPEED TERMINALS

System overhead for terminal output is significant. The total effective asynchronous output rate for a busy 48-user PWB/UNIX system is probably little more than 32 users times 300 baud, or 9,600 baud.

Remember: one 1,200-baud line \approx four 300-baud lines.
one 9,600-baud line \approx thirty-two 300-baud lines.

Be especially wary of high-speed block-mode (CRT) terminals. These can cause loss of input data for other users.

17. LINE PRINTERS

Most line printers are troublesome and impose considerable overhead on the system. Most also lack hardware tabs, character overstrike capability, etc.

18. PAPER-TAPE AND CARD READERS AND PUNCHES

Resist!

19. STAMP OUT UNAUTHORIZED SUPER-USERS

The current PWB/UNIX is not tamper-proof. You can't keep people from "breaking" the system, but you can detect their presence. The following command will list into file */tmp/su* all "set user-id" programs owned by *root* (super-user):

```
find / -user root -a -perm -4100 -a -exec ls -l {} \; >/tmp/su
```

Any surprises in */tmp/su* should be investigated. Related advice:

- Change the super-user password regularly. Don't pick obvious passwords (choose 6-to-8 character nonsense strings that combine alphabetics with digits or special characters).
- If you have dial ports and do not *require* passwords, you are courting trouble.

20. COMMUNICATING WITH YOUR USERS

Many systems use the directory */usr/news* as a billboard of information on specific topics of lasting interest. Users become accustomed to doing "ls -lt /usr/news" to see what has changed recently. Also, it is helpful to have a system mailbox in a known place (*/usr/news/mail*) for announcements of temporary interest; it (as well as the entire */usr/news* directory) should be cleaned out every few weeks, so that nothing older than, say, three months is ever found there.

21. TROUBLESHOOTING

It would be easy to write a book on this topic. The following are some of the key items:

a. Dealing with the hardware service contractor:

- Before you take out a hardware service contract (with DEC* or with someone else), be sure that the contractor agrees to get along with the PWB/UNIX software ("It's the hardware," says you; "It's the software," says the hardware service contractor).
- Keep on top of problems. For instance, DEC has a problem-aging priority scheme. Find out about any such scheme that your contractor may have, and make them prove that it is being followed. Remember that an unreported problem is getting no priority at all. If a problem persists, escalate it up your contractor's local management chain; it may also be effective to complain to your contractor's sales representative.
- If you are serious about service to your users, you should have an extended-period service contract (e.g., 16 hours/day, 6 days/week). Arrange for preventive maintenance (PM), non-critical repair, and add-on installation work to be done before or after prime time.
- If you have a service contract, learn the details. In particular, make sure that PM work is scheduled in advance and that it is completed.
- Ask the hardware service contractor to provide and maintain a "site log." You will have to work on the log, as well.
- Make sure that your hardware vendor (as well as your hardware service contractor, if the two are different) agrees to the presence of non-DEC equipment on your system (even if you have none to start with).
- Keep console sheets. Make sure console error messages are shown to your contractor's Customer Engineers.
- Take core dumps after system crashes and interpret results for Customer Engineers.
- Keep down-time records and make sure that your hardware service contractor knows about them.

b. Dealing with the telephone services vendor:

You are most apt to have telephone problems when you rearrange or add equipment. You may also have occasional central office, trunking, and modem failures:

- Be specific with repair operators: tell them that the trouble involves *data* equipment.
- If your first call fails to get results, ask for the "supervisor" on the second call, and, if necessary, escalate further to get the problem solved.

c. Some obvious problem areas:

- Disk Drives—Over 50% of your problems are likely to be related to the disk subsystem. As mentioned earlier, the way to keep your system *up* is to have a spare disk drive. Remember:
 - Disk PM is important.
 - Make sure that the Customer Engineers who service your hardware see the console error messages produced by PWB/UNIX (and that they understand them).
 - Disk failure can ruin a PWB/UNIX file system. The *only* defense is to make a complete, daily file backup! (See "Protecting User Files" above.)

* DEC is a Registered Trademark of the Digital Equipment Corporation, Maynard, MA.

- Dial Ports—In this area, as well as in the area of synchronous data interfaces, there is room for finger-pointing among all your vendors. Check for obvious things:
 - Is the system in “multi-user” mode?
 - Is the *letctys* file OK?
 - Are any cables loose (*both ends*)?
 - In some telephone offices, trunk-hunting is based on 10-number groups. Hunting *between* such groups can fail independently of anything else.

The possibilities for trouble are many. The “decision table” below attempts to describe some alternatives; it is meant primarily for users of DH11 asynchronous interface devices. If you are unfamiliar with the format, (vertical) Rule 3 reads: “If line rings *and* ring light shows *and* computer does *not* answer *and* switching the modem solves the problem, then it is likely to be a telephone company problem; also, busy out that line.”

Asynchronous Line Problems										
Rules: 1 2 3 4 5 6 7 8 9 0										
<i>Condition:</i>										
Line rings	N	Y	Y	Y	Y	Y	Y	Y	Y	Y
Ring light shows on telephone console	—	N	Y	Y	Y	Y	Y	Y	Y	Y
Computer answers	—	—	N	N	Y	Y	Y	Y	Y	Y
Login message received on terminal	—	—	—	—	N	N	Y	Y	Y	Y
Switching modem solves problem	—	—	Y	N	Y	N	—	—	—	—
User can login	—	—	—	—	—	—	N	N	N	Y
Telephone console shows data received	—	—	—	—	—	—	Y	Y	N	—
Problem affects whole DH11 (up to 16 lines)	—	—	—	—	—	—	Y	N	—	—
<i>Diagnosis and/or Action:</i>										
No problem	—	—	—	—	—	—	—	—	—	X
PDP-11 hardware problem likely	—	—	—	X	—	X	X	—	—	—
Telephone problem likely	X	X	X	—	X	—	—	—	X	—
May be a problem with user's terminal	—	—	—	—	—	—	—	X	—	—
Busy out bad line(s)	X	X	X	X	X	X	X	—	X	—

- Synchronous Ports—High-speed synchronous interface devices are even more trouble than dial equipment. The following is a list of potential trouble spots:
 - Your PWB/UNIX software.
 - Your interface device (e.g., DQS11B).
 - Cable to your modem.
 - Your modem.
 - The communications line.
 - Other modem.
 - Other cable.
 - Other interface device.
 - Other system's software.

Think of the finger-pointing possibilities. The best defense is a good line monitor.

- Power Supply Modules—There are a lot of them, and they do fail, more or less regularly. Hard failure can be detected at the console; voltage drift is tougher. Failure of the FP11 (floating-point unit) power supply is onerous because Customer Engineers are likely to work back from the far end of the “bus,” taking a long time to find the problem.

22. SUGGESTED HARDWARE MODIFICATIONS FOR RJE

Optional DQS11B changes:

- To change the 6-second time-out to 12 seconds: “Replace R4-M8506 module with 47k ohms.”
- To use the DQS11B with an ACU (bit 8 of CSR resets *data terminal ready*): “In module M8506, add E1605 to E1613, and E1614 to BS1 (location E16 on module layout; F16 on prints). On back-plane, delete D3K1 to D4P2, and D4P2 to D3U1; add D3U1 to D3K1, and D2S1 to D4P2.”

23. PHOTOTYPESETTING EQUIPMENT AND SUPPLIES

Read this section if you plan to use the phototypesetting software of PWB/UNIX.

Phototypesetter. The phototypesetter and fonts currently supported by PWB/UNIX are manufactured by:

Graphic Systems, Inc.
Executive Drive
Hudson, NH 03051 (603-889-8550)

The phototypesetter is an on-line C/A/T System 1 with a high-speed turret. The external paper tape reader on the typesetter is not needed, because the typesetter is connected to the PDP-11 CPU via a DR11C.

The following modification (developed by DEC Field Service) should be made to the DR11C (without this modification, the system may crash when the typesetter is powered down): "Add two 390-ohm resistors—from E-18 pin 6 to ground, and from E-18 pin 3 to ground. Put a piece of insulating tubing over the leads so that they do not short out the "etch" runs that they cross."

There are several fonts that are normally used:

<i>Name</i>	<i>Part Number</i>	<i>Troff Name</i>
BT Times Roman	802-016A	R
Times Italic	802-013A	I
Times Bold	802-014A	B
BT PI Font #4	829-021B	S
Geneva Regular	803-032B	G
Geneva Regular Italic	803-033B	GI or J
Geneva Medium	803-034B	GM or M

The first three of the these provide the most-often used (serif) typeface. The last three are used when a sans-serif typeface is desired. The fourth font contains a number of Greek characters and mathematical symbols; see *NROFF/TROFF User's Manual* by J. F. Ossanna. Graphic Systems, Inc. offers a variety of other fonts. For *troff* to be able to use these fonts, corresponding font tables must be built and compiled into the directory *lusr/lib/font*. Fonts for which the source font tables are supplied are:

<i>Name</i>	<i>Troff Name</i>
Boston Condensed	BC
News Condensed	C
Century Schoolbook Expanded	CE
Century Bold Italic	CI or CK
Constant Width Courier	CW
Futura Demibold	FD or UD
Text Greek	GR
Geneva Light	L
Geneva Light Italic	LI
Palatino	PA
Palatino Bold	PB
Palatino Italic	PI
Stymie Bold	SB
Stymie Medium Italic	SI
Stymie Medium	SM

Paper and Chemicals. The phototypesetter "prints" onto photomechanical paper, which can be obtained from a photographic supply house and is specified as:

- Kodak Ektamatic Paper, Grade S, Type 2250,
8 in. × 150 ft., Specification 175 (or equivalent).

Also obtainable from such a supply house are the chemicals for the developing process:

- Kodak Ektamatic A10 Activator (or equivalent).
- Kodak Ektamatic S40 Stabilizer (or equivalent).

These chemicals should be ordered in 9.5-liter (2.5-gallon) containers to work with the circulator.

Developer. A Kodak Ektamatic Processor Model 214K (or equivalent) is used to process the paper from the typesetter. A light-proof box attached to the 214K (to hold the output cassette from the typesetter) is called an "Autofeeder" and can be obtained from:

Peripheral Graphics, Inc.
Andover Industrial Center
York Street
Andover, MA 01810 (617-475-9005)

Circulator and Dryer. A circulator and a paper dryer, as well as a shelf for the dryer can be obtained from:

Mohr Enterprises
8015 North Ridgeway Ave.
Skokie, IL 60076 (312-674-8890)

The needed parts are:

- ME-8 Mohrflow: circulator to increase the usability of the chemicals.
- ME-5 Mohrdry: dryer for the photomechanical paper.
- Dryer Extension: shelf to support the dryer; it connects to the circulator cabinet.

Also obtainable from Mohr Enterprises are cleaners for the developer and circulator. Such cleaning is needed every 4 to 6 weeks, depending on the volume of work:

- R-53 Mohrchem Activator Cleaner Concentrate.
- R-57 Mohrchem Stabilizer Cleaner Concentrate.

Each quart bottle makes 9.5 liters (2.5 gallons) of reusable cleaner to clean the tubing, rollers, and tray of the developer and circulator. Equivalent cleaners can be obtained from other suppliers.

PWB/UNIX Operations Manual*M. E. Pearlman*Bell Laboratories
Piscataway, New Jersey 08854*ABSTRACT*

This manual contains a complete description of console operations, as well as step-by-step instructions for normal operator functions.

The information in this manual was gathered from personal experience, *PWB/UNIX* User's Manual* pages, and Digital Equipment Corporation hardware manuals and microprogram flow diagrams.

This manual is intended to be as general as possible, in order to serve for all PWB/UNIX systems. It is suggested that each location add specific information about:

- a. Hardware configuration.
- b. Telephone line configuration.
- c. Specific logging and record-keeping practices.
- d. Contacts for hardware and software problems.
- e. Site-dependent diagnostic procedures.

* UNIX is a Trademark/Service Mark of the Bell System.



CONTENTS

HARDWARE OPERATION -- PDP 11/45, /70	1
PDP 11/45, /70 CONSOLES 1	
CONSOLE DESCRIPTION AND OPERATION 2	
PDP 11/45, /70 CONSOLE DESCRIPTION 2	
CONSOLE OPERATION -- SWITCH REGISTER 3	
CONSOLE OPERATION -- LIGHT TEST SWITCH 5	
CONSOLE OPERATION -- POWER KEY SWITCH 5	
CONSOLE OPERATION -- STARTING AND STOPPING 6	
CONSOLE OPERATION -- CONTROL SWITCH FUNCTIONS 6	
CONSOLE OPERATION -- ADDRESS SELECT KNOB AND DISPLAY REGISTER 6	
CONSOLE OPERATION -- DATA SELECT KNOB AND DISPLAY REGISTER 7	
CONSOLE OPERATION -- DISPLAYING DATA AT A SPECIFIC ADDRESS 7	
CONSOLE OPERATION -- EXECUTION INDICATOR LIGHTS 7	
CONSOLE OPERATION -- ADDRESS ERROR INDICATOR 8	
CONSOLE OPERATION -- PARITY AND PAR ERR INDICATORS (11/70 only) 8	
CONSOLE OPERATION -- MAPPING INDICATORS (11/70 only) 8	
TAPE DRIVE OPERATION -- TU10 9	
TAPE DRIVE OPERATION -- TU16 10	
DISK DRIVE OPERATION -- RP03 11	
DISK DRIVE OPERATION -- RP04 12	
DISK DRIVE OPERATION -- RP05 13	
DISK DRIVE OPERATION -- RP06 14	
BOOT PROCEDURES	15
BEFORE BOOTING 15	
PDP 11/45 DISK BOOT 15	
PDP 11/70 ROM BOOT 17	
11/45 TAPE BOOT 19	
GENERAL PROCEDURES AFTER BOOTING 20	
OPERATOR COMMANDS	21
MOUNT AND UDMOUNT COMMANDS 21	
DEVNAME (PHYSICAL VERSUS BLOCK) 22	
CHECK AND RECHECK 23	
UNIX SYSTEM RECOVERY PROCEDURES	26
AFTER ALL SYSTEM CRASHES 26	
UNIX SHUTDOWN PROCEDURE	27
TO SHUTDOWN THE SYSTEM 27	
FILE PATCHING PROCEDURE	28
ERRORS OF SEVERE IMPACT 28	
ERRORS OF MODERATE IMPACT 28	
NON-DISASTROUS EXPORTS 29	
UNIX INITIALIZATION PROCEDURE	29
TO OPEN THE SYSTEM TO MULTIPLE USERS 29	
FILESAVE PROCEDURE	30
HOW TO MAKE BACKUPS OF BACKS 30	
TAPE BACKUP PROCEDURE	31
HOW TO MAKE RESTORES OF BACKS 31	
FILE RESTORE PROCEDURE	32
RESTORING OPERATING SYSTEMS 32	
RESTORING THE FILE SYSTEM 34	
CONTACTS	35
LOCAL PROJECTS	36

PWB/UNIX OPERATIONS MANUAL

HARDWARE OPERATION — PDP 11/45, /70

PDP 11/45, /70 CONSOLES



PWB/UNIX OPERATIONS MANUAL

HARDWARE OPERATION — PDP 11/45, /70

CONSOLE DESCRIPTION AND OPERATION

The following documentation is primarily intended to describe the PDP® 11/70 console and its operation. Differences in the PDP 11/45 appear within brackets “[]”. Those cases that are applicable to only one of the two systems are clearly labeled as such.

PDP 11/45, /70 CONSOLE DESCRIPTION

The PDP 11/45 and 70 consoles are composed of the following:

1. Power Key Switch (OFF/POWER/LOCK)
2. Mapping Lights (11/70 only).
3. Seven Execution Indicator Lights displaying the following Central Processor (CPU) States: RUN, PAUSE, MASTER, USER, SUPERVISOR, KERNEL, DATA.
4. A 22-bit [18-bit]¹ Address Register Display.
5. A 16-bit Data Register Display.
6. An Addressing Error (ADRS ERR) Indicator Light.
7. Parity Error (PAR ERR) Indicator Light (11/70 only).
8. High and Low Parity Indicator Lights (11/70 only).
9. A 22-bit [18-bit] Switch Register (CSW).
10. Control Knobs:

Address Display Select which can be set to any of 8 positions:

- a. USER I
- b. USER D
- c. SUPERVISOR I
- d. SUPERVISOR D
- e. KERNEL I
- f. KERNEL D
- g. PROGRAM PHYSICAL
- h. CONSOLE PHYSICAL

Data Display Select which can be set to the following 4 positions:

- a. DATA PATHS
- b. BUS REGISTER
- c. FPP μ ADRS/CPU μ ADRS
- d. DISPLAY REGISTER

11. Lamp Test Switch
12. Control Switches which are labelled:
 - a. LOAD ADRS (load address)
 - b. EXAM (examine)
 - c. CONT (continue)
 - d. ENABL/HALT (enable)
 - e. S-INST/S-BUS CYCLE (single instruction/bus cycle)
 - f. START

1. This documentation shows differences for the PDP 11/45 within brackets “[]”.

PWB/UNIX OPERATIONS MANUAL

HARDWARE OPERATION — PDP 11/45, /70


- g. DEP (deposit)
- h. REG EXAM (register examine, 11/45 only)
- i. REG DEP (register deposit, 11/45 only)

CONSOLE OPERATION — SWITCH REGISTER

The Switch Register consists of 22 [18 on the 11/45] switches labelled 0 through 21[17] (numbers correspond to bit positions). They are used to manually enter both addresses and data into the processor.


To enter an address such as 173020₈, the bits must be divided into groups of three starting from the right. Bits 00-02 in the first group, bits 03-05 in the second group, 06-08 in the third, 09-11 the fourth, 12-14 in the fifth, etc. In each group of 3 bits an octal digit is indicated as follows:

zero




All 3 switches down

one



Lowest numbered switch (right-most bit) up.

two




Middle switch up.

three



Right and middle switches up.

four



Left switch only, in group of three, up.

five



Left and right switches up.

PWB/UNIX OPERATIONS MANUAL


HARDWARE OPERATION — PDP 11/45, /70

six



Left and middle switches up.


seven



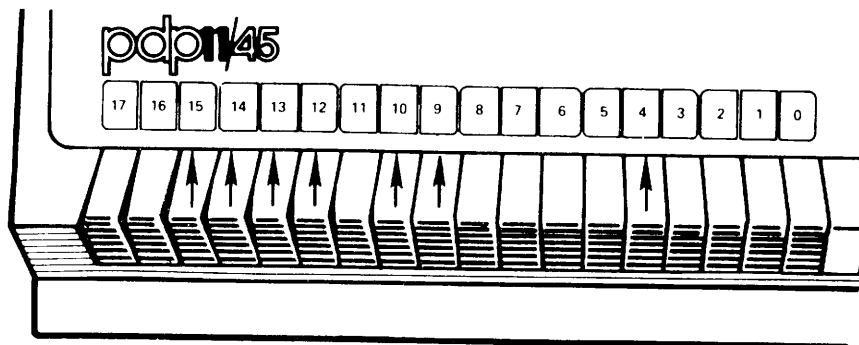
All 3 switches up.

HARDWARE OPERATION — PDP 11/45, /70

The arrows on the following two diagrams indicate which switches should be up to set 173020₈ in both the 11/70 and the 11/45.



one seven three zero two zero



one seven three zero two zero

CONSOLE OPERATION — LIGHT TEST SWITCH

Depressing the Lamp Test switch should light all lamp indicators on the console. Any burnt-out lamps should be replaced at the next preventive maintenance session.

CONSOLE OPERATION — POWER KEY SWITCH

The Power Key controls CPU power in the following manner.

- | | |
|-------|--|
| OFF | Power is off for CPU. |
| POWER | Power is on for CPU. <u>This is the normal position</u> to enable use of all console controls. |
| LOCK | Power is on for CPU. Disables all control switches except for the Switch Register, Data Select, Address Select, and Power Key. |

HARDWARE OPERATION — PDP 11/45, /70

CONSOLE OPERATION — STARTING AND STOPPING

Starting

Once the power is on, execution can be started by setting ENABL/HALT toggle up, putting the starting address in the Switch Register, and depressing the LOAD ADRS Switch. Verify in the Address Register Display lights that the address was entered correctly, then depress START. The computer system RESETS and begins execution of the program instructions at the address specified by the current contents of the Address Display. Depressing start again has no effect, when the CPU is in the RUN state.

When START is depressed with the ENABL/HALT toggle in the halt position (down) the system will RESET, but not commence execution.

Stopping

Depressing the ENABL/HALT toggle will HALT execution.

Continuing

After the computer has been halted, execution can be resumed at the same point it was when halted by depressing the CONT switch. If still in the halted state CONTINUE will single cycle (either by CPU instruction or by Bus² cycle) with each depression (depending on the setting of the S-INST/S-BUS CYCLE switch). If ENABLE is set, normal execution will resume.

CONSOLE OPERATION — CONTROL SWITCH FUNCTIONS

LOAD ADRS loads the contents of the Switch Register into the CPU and displays it in the Address Display. The address displayed in the Address Display depends on the position of the Address Select knob.

The EXAM Switch being depressed causes the contents of the location specified in the Address Display to be displayed in the Data Display lights when the Data Select knob is in the DATA PATHS position.

The DEP switch being raised causes the current contents of the Switch Register to be deposited into the address specified by the current contents of the Address display.

CONT being depressed after the system is halted, causes the CPU to resume execution as described in the previous section on STARTING AND STOPPING.

ENABL/HALT is a two position switch whose functions also have been described in the previous section.

S/INST-S/BUS affects only the operation of the CONTINUE switch. It controls whether execution stops after a single INSTRUCTION or a single BUS CYCLE. This switch has no effect when CPU is in the RUN state.

START starts execution or RESETS depending on the setting of the ENABL/HALT switch, as previously described.

REG EXAM (11/45 only) depressed causes the contents of the General Purpose Register specified by the low order 4 bits of the Bus Address Register to be displayed in the Data Display lights.

REG DEP (11/45 only) raised causes the contents of the Switch Register to be deposited into the General Purpose Register specified by the current contents of the CPU Bus Address Register.

FOR FURTHER INSTRUCTIONS SEE THE PDP 11/45 PROCESSOR HANDBOOK.

CONSOLE OPERATION — ADDRESS SELECT KNOB AND DISPLAY REGISTER

Address Select provides an interpretation for the Address Display Register when set as follows:³

2 The Bus (or UNIBUS) is the primary control and communications path connecting most of the PDP 11 system's components and peripherals.

HARDWARE OPERATION — PDP 11/45, /70

VIRTUAL	(6 positions for USER, SUPERVISOR, and KERNEL) Indicates the current address as a 16-bit Virtual address (when the Memory Management Unit is turned on, i.e. UNIX is running, otherwise it indicates the true 16-bit Physical Address). <u>FOR DEBUGGING ONLY</u>
PROGRAM PHYSICAL	Displays a true 22-bit [18-bit] Physical Address for the current address reference. <u>ALSO FOR DEBUGGING</u>
CONSOLE PHYSICAL	Display a 22-bit [16-bit] Physical Address to be used for ALL console operations such as LOAD ADRS, EXAM and DEP.

CONSOLE OPERATION — DATA SELECT KNOB AND DISPLAY REGISTER

The Data Select Knob controls the contents of the 16-bit Data Display Register according to the following settings:

DATA PATHS	The <u>normal mode</u> when performing console operations. Shows examined or deposited data.
DISPLAY REGISTER	Displays the current contents of the 16-bit write only "Switch Register". This is the normal position for the Data Select Switch when UNIX is running.
BUS REG	Displays the contents of the Bus Register.
μ ADRS FPP/CPU	Bits 15-08 display the current ROM ⁴ address of the Floating Point Processor microprogram. Bits 07-00 display the current ROM address of the CPU microprogram.

CONSOLE OPERATION — DISPLAYING DATA AT A SPECIFIC ADDRESS

To display the data contained at address X_8 :

1. Set Power Key to POWER position.
2. Set Address Select Knob to CONSOLE PHYSICAL.
3. Set Data Select Knob to DATA PATHS.
4. Enter X_8 into the Switch Register.
5. Depress LOAD ADRS.
6. Depress EXAM.
7. The Data Display lights will contain the data from address X_8 .
8. Depressing EXAM again will display the data from address $X+2_8$.

CONSOLE OPERATION — EXECUTION INDICATOR LIGHTS

Lighting of each display indicates the following:

RUN	The processor is executing instructions, including WAIT instructions, but is not in a Pause cycle.
-----	--

3. This documentation shows differences for the PDP 11/45 within brackets "[]".

4. ROM is an abbreviation for "read only memory".

PWB/UNIX OPERATIONS MANUAL

HARDWARE OPERATION — PDP 11/45, /70

PAUSE	The processor is in a Bus Pause or Interrupt Pause Cycle, waiting for a Unibus device or Memory. Or, the CPU has been halted from the console.
MASTER	The CPU is in control of the Unibus or during Console operations.
KERNEL SUPER USER	These lights show where the current memory reference is taking place.
DATA	Shows whether I (program) or D (data) space is being used in the current memory reference. It is on when D space is used and off when I space is used.

CONSOLE OPERATION — ADDRESS ERROR INDICATOR

ADRS ERR is on when an addressing error occurs. (e.g., non-existent memory, access control violation, reference of unassigned memory pages).

CONSOLE OPERATION — PARITY AND PAR ERR INDICATORS (11/70 only)

PARITY indicators display the parity bits associated with the HIGH and LOW bytes of the word currently read from Cache memory. Indicators are off during write operation.

PAR ERR indicator is lit when a Unibus or a memory parity error is flagged.

CONSOLE OPERATION — MAPPING INDICATORS (11/70 only)

The 16, 18, and 22 indicate the bits used for Memory Management mapping for each cycle.

PWB/UNIX OPERATIONS MANUAL

HARDWARE OPERATION — PDP 11/45, /70

TAPE DRIVE OPERATION — TU10

Please file Chapter 3 of DEC® MAINTENANCE MANUAL DEC-00-TU10S-DC here.

PWB/UNIX OPERATIONS MANUAL

HARDWARE OPERATION — PDP 11/45, /70

TAPE DRIVE OPERATION — TU16

Please file Sections 1.7 and 1.8 of DEC MAINTENANCE MANUAL EK-TU16-MM-001 here.

PWB/UNIX OPERATIONS MANUAL

HARDWARE OPERATION — PDP 11/45, /70

DISK DRIVE OPERATION — RP03

Please file Section 3 of DEC ISS MANUAL UD002341-1 here.

PWB/UNIX OPERATIONS MANUAL

HARDWARE OPERATION — PDP 11/45, /70

DISK DRIVE OPERATION — RP04

Please file Section 3 of DEC ISS MANUAL UD002511-1 here.

PWB/UNIX OPERATIONS MANUAL

HARDWARE OPERATION — PDP 11/45, /70

DISK DRIVE OPERATION — RP05

Please file “operation” section of DEC ISS MANUAL for RP05 disk drive here.

PWB/UNIX OPERATIONS MANUAL

HARDWARE OPERATION — PDP 11/45, /70

DISK DRIVE OPERATION — RP06

Please file "operation" section of DEC ISS MANUAL for RP06 disk drive here.

BOOT PROCEDURES

BEFORE BOOTING

Make sure the system is *idle*.⁵ If *UNIX* is running, the system should be cleared via the UNIX SHUT-DOWN PROCEDURE.


The normal boot performed is a ROM BOOT. The TAPE BOOT is performed (on the 11/45 only) when loading a special system, or when a ROM BOOT is impossible.

When performing any boot procedure the Power Key Switch should be in the *POWER* position, the Address Select Knob should be in the *CONSOLE PHYSICAL* position, and the Data Select Knob should be in the *DATA PATHS* position.

Before performing any procedure, please read that procedure completely.

PDP 11/45 DISK BOOT

- A. Halt system (set *ENABL/HALT* toggle down).
- B. Depress *START* (this resets the system).
- C. Set *CONSOLE ADDRESS SWITCHES (CSW)* to the address of your ROM or Boot Strap Loader (e.g., 173020).⁶




5. The system is *idle* when no *user* processes are running.

6. See *romboot(VIII)* in the PWB/UNIX User's Manual for further details.

BOOT PROCEDURES

- D. Depress *LOAD ADDRESS* toggle.
- E. Set *ENABL/HALT* toggle up.
- F. Set console switches to appropriate settings.
 1. For *SINGLE-USER*, set 173030₈ → *CSW*.



2. For *MULTI-USER*, leave *CSW* at original settings.

- G. Depress *START*.
- H. At console, enter:⁷

00 = unix©

You enter the two zeroes, the system name,⁸ and the ©, all on the same line as the #.

7 Information or commands entered by the operator appear underlined. The symbol © is used throughout this procedure to indicate a carriage return.


8 If you are booting a special system, the name of that special system may be substituted for "unix". See *diskboot*(VIII) for complete boot instructions.

PWB/UNIX OPERATIONS MANUAL

BOOT PROCEDURES

PDP 11/70 ROM BOOT

- A. Halt system (set *ENABL/HALT* toggle down).
- B. Depress *START* (this resets the system).
- C. Set *CONSOLE ADDRESS SWITCHES (CSW)* to 165000_8 .



- D. Depress *LOAD ADDRESS* toggle.
- E. Set *ENABL/HALT* toggle up.
- F. Set *CSW* to 000000_8 (all *CONSOLE ADDRESS SWITCHES* down).⁹
- G. Depress *START*.
- H. Wait for console to respond with a #.


9. If the default boot address on your ROM is not set to the correct boot device, entering zeroes will fail. See *11/70boot(VIII)* for complete instructions.

PWB/UNIX OPERATIONS MANUAL

BOOT PROCEDURES

I. Set console switches to appropriate settings:

1. For *SINGLE-USER*, set $173030_8 \rightarrow CSW$.



2. For *MULTI-USER*, leave console switches at original settings.

J. At console, enter:

00 = unix©


You enter the two zeroes, the system name,¹⁰ and the ©, all on the same line as the #.

¹⁰ If you are booting a special system, the name of that special system may be substituted for "unix". See *diskboot*(VIII) for complete boot instructions.

BOOT PROCEDURES

11/45 TAPE BOOT

- A. After the system is halted (*ENABL/HALT* toggle down), depress *START* (this resets the system).
- B. Mount the appropriate tape on the tape drive (make sure tape is ready at *load point*).
- C. Set the address of your ROM or Boot Strap Loader in the console switches (e.g., 173000₈).



- D. Depress *LOAD ADDRESS* toggle.
- E. Set *ENABL/HALT* toggle up.
- F. Set console switches to appropriate settings:
 - 1. For *SINGLE-USER*, set 173030₈ → *CSW*.
 - 2. For *MULTI-USER*, leave console switches at original settings.
- G. Depress *START* (system will give you =).
- H. At console, enter:
 - 1. If loading *UNIX*
`unix©`
 - 2. If loading a special system
`"filename"©`
 (where *filename* is the name of the system being loaded, e.g., "pwbx0409")

PWB/UNIX OPERATIONS MANUAL

BOOT PROCEDURES

GENERAL PROCEDURES AFTER BOOTING

After the boot is completed, UNIX will print a "restricted rights" message and tell you the System Name and the amount of memory available. The latter should be checked against a previous boot to be sure that all of memory is available. Otherwise, serious problems can occur.

- A. If booting as part of *SYSTEM RECOVERY*, continue *SYSTEM RECOVERY PROCEDURE*.
- B. If bringing up *UNIX MULTI-USER*, perform *UNIX INITIALIZATION PROCEDURE*.

OPERATOR COMMANDS

MOUNT AND UMount COMMANDS

The **mount** and **umount** commands are used to logically **mount** or **unmount** filesystems. The **mount** command without any arguments is used to list which filesystems are *mounted* and on which *device* they are located. This aspect of the command will not give accurate data when performed on a SINGLE-USER system.

The format for these commands is:

mount

Output will be formatted:

“*fsname on devname permissions on day date time year*”

An example would be:

“/a3 on /dev/rp5 read/write on thu may 15 17:43:51 1976”

or

mount devname fsname option¹¹

or

umount devname

devname = the logical device name (e.g., “/dev/rp12”)

fsname = the filesystem name that you wish this device to be called (e.g., “/f3” or “/bck”)

permissions = *read only* or *read/write* permissions are granted.

option = the optional “-r” for *read only* (this field is left blank for write capabilities)

Before **mounting** any filesystems, it is a good practice to perform a **mount**, without arguments, to know which devices you will have to re-mount, and what they must be named.

If the **umount** command complains that the device is *busy*, it means that someone is using the filesystem (maybe you). Change directory (**chdir** (1)) to “/” and try again.

11. NOTE: Before **mounting** a filesystem, be sure that a filesystem with the same name (i.e., *fsname*) is not already mounted.

PWB/UNIX OPERATIONS MANUAL

OPERATOR COMMANDS

DEVNAME (PHYSICAL VERSUS BLOCK)

Block device names are used in commands which must have data passed through the system buffers (e.g., **mount**, **umount**). They are formatted:

"/dev/rpl2" or "/dev/mt0"

Physical device names should always be used for *file maintenance* commands (e.g., **vc10**, **vc88**, **vc50** and **check**). Use of Physical device names makes the transfer of data "device-to-device", thereby shortening execution time considerably. These are formatted:

"/dev/rrp12" or "/dev/rmt0"

(note the extra "r" in "rrp12")

Physical device names should be used only in referencing logically **umounted** filesystems, since changes made to the physical device may not be reflected in the systems in-core image of the device. The exception to this rule is the "root" device, which is never **umounted**.

OPERATOR COMMANDS

CHECK AND ICHECK

The **check** command is used to check filesystems for possible file damage. For example, **check** should always be used after a system crash.

The **check** and **icheck** commands are formatted:

```
# check option devname devname ...
```

and

```
# icheck option devname
```

Leaving the *devname* field blank for **check** will default to a **check** of all filesystems listed in the file "/etc/checklist".

```
option =          option field.12
```

The **-s** (for salvage) is used with **icheck** to create a new *free list* for the filesystem specified; **icheck -sX** (where "X" is 3 for RP03s, and 4 for RP04s and RP05s) should always be performed on an **unmounted** filesystem (except for the root which is always mounted); **icheck -sX** of the root must be followed by an IMMEDIATE reboot.

```
devname =        physical device name
```

For an example, the *root* filesystem for a system might be located on /dev/rrp0. If you wanted to check only that particular filesystem you would type:

```
# check /dev/rrp0
```

On the other hand:

```
# check
```

would check all of them.

An example of what the normal output of check should look like is:

```
# check /dev/rrp0
```

```
/dev/rrp0:
spcl      155
files     253
large     85
direc     18
indir     85
used      2849
last      6681
free      3698
```

¹² An explanation of all possible options can be found in the description of *check* (VIII) and *icheck* (VIII)

PWB/UNIX OPERATIONS MANUAL

OPERATOR COMMANDS

The two examples below illustrate errors of moderate impact and also non-disastrous errors called *100s* and *20/s*. They are not particularly dangerous, and the system can run with them present, but they should be repaired before opening the system to users.

```
# check /dev/rrp0
```

```
/dev/rrp0:  
 372 100  
spcl      155  
files     253  
large     85  
direc     18  
indir     85  
used     2849  
last     6681  
free     3698
```

```
# check /dev/rrp0
```

```
/dev/rrp0:  
 372 100  
 465 201  
 379 201  
spcl      155  
files     253  
large     85  
direc     18  
indir     85  
used     2849  
last     6681  
free     3698
```

The following errors should be repaired before users log on via **icheck**; they indicate errors in the *free list*. These types of errors leave the system virtually useless.

```
# check /dev/rrp0
```

```
/dev/rrp0:  
Bad freeblock  
15 dups in free  
2005 missing  
spcl      155  
files     253  
large     85  
direc     18  
indir     85  
used     2849  
last     6681  
free     3698
```


OPERATOR COMMANDS

These errors illustrate *Errors of Severe Impact* (377s and 177s) and if encountered, they should be repaired immediately via **patchup**. It is disastrous to run with such errors outstanding, since these types of errors will spread throughout the filesystem.

```
# check /dev/rrp0
```

```
/dev/rrp0:  
 675 377  
 496 177  
spcl      155  
files     253  
large     85  
direc     18  
indir     85  
used      2849  
last      6681  
free      3698
```

There are many other types of *serious* errors that can appear in your **check** output. The repair of these errors is too complicated to be described in this document.¹³

Phone numbers to call when problems arise can be found in the CONTACTS section of this book.

¹³ For more details on repairing file system errors, see *Repairing Damaged PWB/UNIX File Systems* by P. D. Wandzilak.

UNIX SYSTEM RECOVERY PROCEDURES

AFTER ALL SYSTEM CRASHES

- A. Note all error indications, hardware or software:
for example;
 1. CPU looping or halted.
 2. Power on?
 3. Error messages on console.
 4. RP04 *DISK UNSAFE*.
- B. Shutdown the System to all users.
 1. Halt System, by pressing *halt* switch down.
 2. Busy out (push buttons in) dial up lines for that system (check labels on *data-set* consoles).
- C. Attempt System Core Dump Procedure.
 1. Make sure the *KEY* is in the *POWER* position (not *LOCK*)
 2. Mount dump tape with write-ring on the tape drive. Make sure the tape is *ONLINE* and ready at *load point*.
 3. Depress *START*.
 4. Set $044_8 \rightarrow CSW$.
 5. Depress *LOAD ADDRESS* switch down.
 6. Set *ENABL/HALT* switch up.
 7. Depress *START* (dump will be written to tape).
 8. When dump tape has rewound, unmount dump tape; remove write-ring; and label the tape with the following information:
 - a) "CORE DUMP"
 - b) date
 - c) time
 - d) System
 - e) Console messages (e.g., KA6 #, APS #, and trap type).
- D. Reboot system *SINGLE-USER*, via ROM BOOT PROCEDURE.
- E. Do a *CHECK* of the filesystems, via **check**.
check
- F. Perform FILE PATCHING PROCEDURE.
- G. Reboot the system for 'MULTI-USER' (via either the *ROM BOOT* Procedure or changing the console switches to any other value other than the previous setting, followed by typing a "CTRL/d").
- H. Bring up *UNIX* via UNIX INITIALIZATION PROCEDURE.
- I. Record downtime in SYSTEM DOWN LOG.

UNIX SHUTDOWN PROCEDURE

TO SHUTDOWN THE SYSTEM

- A. Make sure system is *idle* (all users logged-off). This can be accomplished by executing the **shutdown** program.¹⁴

```
login: root
password: !;%&?
# shutdown
```

If the shutdown program fails, the following steps can be used instead.

1. **BROADCAST** a message until all users, who can, are logged-off.

```
# wall FILENAME
(to broadcast the contents of a file)
```

OR

```
# wall
The body of you message should be typed here.
"CTRL/d"
```

2. Busy out (push in buttons) the dial-up lines for the appropriate system(s).
3. Halt *HASP* or *HASP2*, if applicable.

```
# setuid hasp
% /usr/hasp/hasphalt
# setuid hasp2
% /usr/hasp2/hasp2halt
(Systems will respond with either hasp halted or hasp2 halted or both)
```

4. On systems with UNIVAC RJE.

```
# setuid uvac
% /usr/uvac/uvachalt
(system will respond with uvac halted)
```

5. Do a **ps a** to find out the process number of the *CRON*, then, **kill** the *CRON*, if it is active, and any other active processes other than the **ps a** and your terminal process.

```
# ps a
(system will give you a list of active processes)
# kill -9 process#
(where process# is the decimal process number of each active process)
```

6. Enter

```
# sync;sync;sync
```

- B. Halt system (set *HALT* toggle down).

- C. If turning off power, make sure the disk drives have been stopped before turning **CONSOLE POWER KEY** to the OFF position.

¹⁴ Commands actually entered by the operator appear underlined.

FILE PATCHING PROCEDURE

After a system crash, there are several types of damage that can disable a file. Some of these errors must be repaired immediately before the disease spreads throughout the entire filesystem. A complete analysis of these errors can be found in the explanation of **check** in the COMMANDS section of this manual.

ERRORS OF SEVERE IMPACT

Errors such as *377s*, *177s*, and those with a format of:

*nnn complaint; i = xxxx (class)*¹⁵

should be treated immediately. Errors in the *root* should be handled first, because the commands *you* execute are in the *root*.

377s and *177s* can be handled by the program **patchup**.

patchup

The *complaint* types of errors are too complicated to have established handling procedures, so an immediate telephone call to SYSTEM SUPPORT is recommended.

ERRORS OF MODERATE IMPACT

Errors of the following formats are often found after a crash, and indicate problems in the *super block*:

n dups in free

bad freeblock

n missing

After all other errors are repaired, these must be remedied by:

icheck -sX devname (physical device name)

where X is 3 for RP03s and 4 for RP04s and RP05s.

icheck -sX must always be performed on an unmounted filesystem, except in the special case of the *root* device, which is always mounted. Errors in the *super block* of *root* **MUST** be repaired first and be followed by an IMMEDIATE reboot.

NON-DISASTROUS ERRORS

201s and *100s* are fairly innocuous errors. They can be repaired by the **patchup** program as described for *377s* above.

¹⁵ A full explanation of this type of error can be found in the description of *check* (VIII).

PWB/UNIX OPERATIONS MANUAL

UNIX INITIALIZATION PROCEDURE

TO OPEN THE SYSTEM TO MULTIPLE USERS

A. If running *SINGLE-USER*:

1. Do a **CHECK** of the filesystems via **check** (if the filesystems are not clean, refer to FILE PATCHING PROCEDURE).

check

2. Set date to the correct time.
3. Reboot by changing the position of the *CSW* to any position other than the original setting, and entering *CTRL/d*.¹⁶

B. When running *MULTI-USER* log on as *SUPER USER*.

- login: root

- password: !;%&? (system will give you a #)

If you are running *SINGLE-USER* you automatically get back the #.

C. ONLY IF the filesystems are clean, to open the system for general use (i.e., to allow users to log on):

1. Start */ETC/STARTUP*.

startup

This command alone will activate the processes for each of the systems. The file *"/etc/startup"* should contain any commands and cleanup to be done on each reboot.

2. Do a **ps a** to make sure *HASP*, *CRON*, *HASP2*, and *UVAC* (if applicable) are active.

ps a

3. Unbusy the appropriate dial-up lines for each system.
4. Log-off (*CTRL/d*).

16. Do not do a *CTRL/d* after an *icheck -sX* of the *ROOT* device; instead the system MUST be halted and completely rebooted.

PWB/UNIX OPERATIONS MANUAL

FILESAVE PROCEDURES

HOW TO MAKE BACKUP PACKS

- A. Before initiating **filesave**, check the *mail* for the *root* to make sure there are no files to be restored. Restore any files at this time, because when you execute the **filesave** you may be destroying the only good backup you have for a file. For instructions on how to restore a file, see the FILE RESTORE PROCEDURE section of this manual.
- B. Make sure the system is *idle* (no user processes active). This can be accomplished by executing the **shutdown** program.¹⁷

login: root

password: !;%&?

shutdown

If the shutdown program fails, an explanation of the steps that can replace it can be found in the UNIX SHUTDOWN PROCEDURE section of this book.

- C. Halt the system (set *HALT* toggle down).
- D. Reboot the system *SINGLE-USER* via ROM BOOT PROCEDURE.
- E. Perform file maintenance.
1. Make any necessary patches to the filesystems (see FILE PATCHING PROCEDURE).
 2. Get packs and tapes for the File Save from the tape/disk cabinets using the oldest dated disk packs first. To determine the appropriate disk and tape, check your Disk Log and Tape Log.
 3. Start **FILESAVE** *shell* procedure.

filesave.n

(where *n* is the system identification letter, e.g., *a*)

Allow the **filesave** to run to completion. If any part of the file save fails, re-execute starting from that point.

4. Record all disk packs and tapes used on the log sheets.
- F. Bring up UNIX MULTI-USER via the UNIX INITIALIZATION PROCEDURE.

¹⁷ Commands actually entered appear underlined

TAPE BACKUP PROCEDURES

In addition to the backup packs created in the **filesave**, tape backups for each filesystem should be created once a week. One filesystem fits on each tape used, provided you are using your tape media at 1600 bpi. These tapes serve as long-term storage for files.

The best way to make a backup tape is to make a copy of a Backup disk. This way tapes can be made while the system is available to the users.

HOW TO MAKE A BACKUP TAPE

- A. A backup pack for the filesystem of which you wish to make a tape copy must be physically mounted on the spare drive. If no spare drive is available, tapes must be made directly from the filesystem that is normally on-line, on a *single-user* system. This can be done immediately after the **filesave**.
- B. Tapes and packs for each backup should be found ahead of time. Tapes should not be reused before approximately 6-8 weeks. The correct backup pack and tape to use can be determined by checking the Disk and Tape Log(s).
- C. To create a backup tape for a filesystem a **vc10** is performed (see **volcopy** (VIII)).¹⁸ The type of device you copy to depends on the type of tape drive used:

For a 800 bpi(TU10) tape drive, use:

```
vc10 fsname /dev/rrpX n packid /dev/rmt0 tapeid
```

For a 1600 bpi(TU16) tape drive, use:

```
vc10 fsname /dev/rrpX n packid /dev/rmt1 tapeid
```

- *fsname* is the filesystem name (e.g., *al* or *root*).
- *X* is the number of the disk drive on which the *fsname* or its backup is mounted.
- *n* is the number of the section of the pack on which the filesystem resides (e.g., 0, 1, 2, 3, 4, 5, or 6).

Examples:

/dev/rrp14	section 4 on drive 1
/dev/rrp41	section 1 on drive 4
/dev/rrp0	section 0 on drive 0
/dev/rrp65	section 5 on drive 6
/dev/rrp64	section 4 on drive 6

- *packid* is the five-character external pack label (e.g., p0452).
- *tapeid* is the five-character external tape label (e.g., t0352).

Example:

```
vc10 root /dev/rrp0 p0450 /dev/rmt1 t0432
```

- D. After backup tapes are created, packs should be returned to original locations and the tape identification numbers should be recorded in your Tape Log.

¹⁸ During the creation of a backup tape errors may occur. If "read errors" occur, the reliability of the disk pack is questionable. In any case the backup tape will be bad. Detailed description of these errors may be found below under RESTORING FILES FROM TAPE.

FILE RESTORE PROCEDURE

During a system crash, files may be lost. Users may also accidentally destroy their own files. In either event, the user will probably call or mail to the *root* to request that his/her file(s) be restored. When the user calls, make sure you get the complete *pathname*, the time and date of the desired restoral source, and the approximate size of the file. *Pathnames* are composed as follows:

fsname/dirname/filename

fsname = the *filesystem* name (e.g., *b3*).

dirname = the complete directory name, including all sub-directories between the file and the *fsname* (e.g., *source/s2* or *actg*).

filename = the actual file name.

The operator should perform any outstanding file restores before all filesaves that could destroy the only good restoral source.

There are, at present, two basic methods of restoring files: (1) from tape, or (2) from disk. The latter method is preferred, if you have a pack that contains a copy of the file.

RESTORING FILES FROM DISK

- A. Make the file accessible on an available spare disk drive, by performing the following steps:
1. Carefully, physically mount the backup pack (found by checking your Disk Log) on the drive and bring it to *READY*.
 2. **Mount**, as *read only*, the device on which the backup filesystem resides. In the **mount** command *devname* is the block device name of the filesystem being mounted (e.g., *ldev/rp12*). The filesystem should be **mounted** as */bck* or some other currently unused mount point directory.

```
# mount devname /bck -r
```

3. Change your working directory to *"/bck/dirname"* (explanation of *dirname* is the same as above).

```
# chdir /bck/dirname
```

- B. Check to make sure a copy of the file exists.

```
# ls -l filename
```

The date the file was last written to and the number of characters it contains are pertinent information in determining that particular validity of the copy.

- C. Copy the file into the desired *pathname* (the pathname given by the user). This may be accomplished as follows:

```
# cp filename fsname/dirname/filename
```

- D. Do a **ls -l** of the copied file.

```
# ls -l filename
```

Check this against the original **ls -l** of the file. If the number of characters agree, the copy was successful. Check for the owner and the modes of the original file.

PWB/UNIX OPERATIONS MANUAL

FILE RESTORE PROCEDURE

- E. Notify the user either through the UNIX mail command or by telephone.
- F. Return the system to its original configuration:
 1. Logically **umount** the filesystem copy (*devname* for **umount** is also block, e.g., */dev/rp12*).
 - # chdir /
 - # umount *devname* (watch for error messages, if any occur restart from preceding step)
 - # mount (to make sure the spare filesystem is no longer mounted)
 2. Log-off (*CTRL/d*)
 3. Remove the pack using the following steps:
 - a) Turn off the disk drive containing the unmounted filesystem, by pressing the button labelled *STOP*.
 - b) Wait until the pack has completely stopped spinning.
 - c) Carefully remove the disk pack from the drive.

PWB/UNIX OPERATIONS MANUAL

FILE RESTORE PROCEDURE

RESTORING FILES FROM TAPE

A. Copy the file from tape onto a disk.

- 1 Find an available device name on disk, where the data from tape can be read and stored (i.e. a *devname* that is logically **unmounted** where there is no valuable data). This usually involves physically mounting a scratch disk pack on a spare drive.
- 2 Find the appropriate tape determined by checking the date on the tape labels.
- 3 Mount the tape on the tape drive on the system where the *devname* is available.
- 4 **Login** as *root* and **read** the tape onto the available *devname* (physical *devname*, e.g., */dev/rrp10*).

login: root

password: !;%&?

vc10 *fsname* /dev/rmt0 txxxx *devname* pxxxx (for 800 bpi tapes)

or

vc10 *fsname* /dev/rmt1 txxxx *devname* pxxxx (for 1600 bpi tapes)

Messages like *read error block N*¹⁹ can *not* be ignored, and watch for fatal errors such as *read error -1*; **vc10** should say *END* and tell you the number of blocks copied, when it is finished.

- #### B. Now you have a copy on disk of the filesystem in which your file resides. You can now perform steps *A2* through *F2* of RESTORING FILES FROM DISK (without physically removing any packs that are already mounted).

¹⁹ When this error occurs on the input device, the 10-block record currently being copied is zeroed out (because the error is probably caused by a bad block on the tape). The probability is small, but this may occur in the very file you are attempting to restore, making this tape useless for backup.

PWB/UNIX OPERATIONS MANUAL

CONTACTS


This section should contain the names, addresses, telephone numbers, etc., of your installation Customer Engineers, PWB/UNIX system experts, etc.

PWB/UNIX OPERATIONS MANUAL

LOCAL PROCEDURES

This section is reserved for your installation's local procedures.







Repairing Damaged PWB/UNIX File Systems

P. D. Wandzilak

March 1978

**Repairing Damaged PWB/UNIX
File Systems**

CONTENTS

1	INTRODUCTION	1
2	DEVICES—SPECIAL FILES	1
3	CHECK DIAGNOSTICS	1
3.1	Free-List Diagnostics	2
3.2	Block Diagnostics	3
3.3	Inode Diagnostics	4
3.3.1	Type-100 Error.	5
3.3.2	Type-201 Error.	5
3.3.3	Type-377 Error.	5
3.3.4	Type-177 Error.	6
4	REPAIRING FILE SYSTEM DAMAGE	6
4.1	Free-List Diagnostics	6
4.1.1	Root File System.	6
4.1.2	Non-root File System.	8
4.2	Block Diagnostics	8
4.2.1	“DUP” Diagnostic.	9
4.2.2	An Example of a More Complex “DUP” Diagnostic.	10
4.2.3	“BAD” Diagnostic.	12
4.2.4	“DIN” Diagnostic.	13
4.3	Inode Diagnostics	17
4.3.1	Type-100 Error.	17
4.3.2	Type-377 Error.	18
4.3.3	Type-177 Error.	19
4.3.4	Type-201 Error.	20
5	ACKNOWLEDGEMENTS	22
6	REFERENCES	22

Repairing Damaged PWB/UNIX File Systems

P. D. Wandzilak

Bell Laboratories
Piscataway, New Jersey 08854

1. INTRODUCTION

This document is intended to help PWB/UNIX* operations personnel repair file system damage. It assumes the reader has an understanding of [1,2]. A good understanding of the PWB/UNIX file system structure is also a prerequisite; in particular, it is assumed that the reader is familiar with the concept of *isize*, *fsize*, *ilist*, *inode*, *i-number*, and *small* and *large* files.

The task of repairing damaged file systems is not an easy or straightforward process. A considerable amount of thought and analysis of the problem must first be applied before any surgery is performed. Furthermore, situations do occur when a definite diagnosis of the problem is impossible; in such a situation, only experience and profound courage count.

This document is meant to be introductory in nature, and will not cover cases that exhibit a high degree of complexity or that require extensive knowledge and experience.

2. DEVICES—SPECIAL FILES

The PWB/UNIX Operating System distinguishes between two classes of devices: character devices and block devices. *Character devices* are basically byte oriented, e.g., terminals. *Block devices* are oriented to transferring larger groups of data, e.g., disks, magnetic tapes, etc. Each I/O device supported by PWB/UNIX is assigned a device class and major and minor device numbers. The major number is associated with the software driver for that device. The minor number specifies both the physical drive unit and a designated logical portion of the device. A *special file* is a PWB/UNIX file that has its major and minor device numbers recorded in `addr[0]` of the *inode's* control structure. Examples:

<i>Character Devices:</i>	<i>Block Devices:</i>
/dev/tty8	/dev/rp54
/dev/rrp24	/dev/rp24
/dev/rrp0	/dev/rp0

In the above examples, the devices `/dev/rp0` and `/dev/rrp0` both refer to the same physical device, but they have different access protocols.

3. CHECK DIAGNOSTICS

The first step in repairing file system damage is to find the extent of the damage. This is accomplished by executing the *check* command. Before the *check* command is invoked, the PWB/UNIX system should be rebooted and contained within the single-user environment; *check(VIII)* examines a file system and reports all inconsistencies. The normal output of *check* consists of the file-system device name followed by eight lines of summary information:

- the number of special files;
- the total number of files;
- the number of large files;
- the number of directories;
- the number of indirect blocks;
- the number of blocks used;
- the number of the highest block used;
- the number of free blocks.

The following is an example of the normal *check* output:

* UNIX is a Trademark of Bell Laboratories.

```
# check /dev/rp0
/dev/rp0:
spcl      158
files     261
large     90
direc     29
indir     92
used      3630
last      6684
free      2923
```

Any additional information not shown above represents diagnostic messages produced by *check*. Diagnostic messages follow the file-system device name and precede the summary information, as shown below:

```
# check /dev/rp0
/dev/rp0:
  diagnostic messages:
  :
spcl      158
files     261
large     90
direc     29
indir     92
used      3630
last      6684
free      2923
```

Each diagnostic message is a single entity. The user must have a thorough understanding of these messages in order to diagnose the problem and take proper corrective action.

The problem with most of the diagnostic messages is that their explanation is implicit rather than explicit. The following discussion clarifies these messages and advises the reader about their significance and importance.

3.1 Free-List Diagnostics

These diagnostics are the most common of all the *check* diagnostic messages. They pertain to areas within a file system that are designated as *free*. Each free area is linked together to form a chain, called the *free list*. The free list is dynamic and therefore sensitive to abrupt or abnormal system stops. When such a stop does occur, the free list can become curdled; this happens when the head of the in-core *free list* is not written out to disk. The *check* command recognizes these conditions and reports them via various diagnostic messages. The message "bad freeblock" means that a block number in the free list resides outside the range of available space. The message "dups in free" indicates duplication of disk blocks; this occurs when the same block or set of blocks in the free list appears in other files or in an earlier part of the free list. The last message type in this category concerns "missing" block numbers. A block is considered missing when it resides neither in the *ilist* (which contains control blocks for all files in a file system) nor in the remaining portion of the file system that consists of directories, ordinary files, and the free list.

When the *check* command is invoked upon a file system, and if the file system's free list is not consistent, one or more of these three diagnostic messages will appear. The output of *check* for a file system containing these types of free list errors is shown below:

```
# check /dev/rp0
/dev/rp0:
bad freeblock
20 dups in free
2000 missing
spcl      158
files     261
large     90
direc     29
indir     92
used      3630
last      6684
free      923
```

3.2 Block Diagnostics

This section describes diagnostic messages that concern disk blocks of both ordinary files and directories. An *inode* is a 32-byte control block structure that identifies various attributes of each file or directory (see *File System(V)*). All inodes reside in the portion of the file system identified as the *ilist*. Each *inode* in the *ilist* is numbered and identified by an i-number. An *i-number* uniquely identifies a file in a file system. This type of message is of the form:

b# error; inode = i# (block type)

where *b#* indicates the actual *decimal* block number, *error* is the type of diagnostic in a mnemonic form, and *i#* represents the actual i-number associated with the particular block number.

These diagnostics are based upon information contained in the super-block of the file system; although limited precautions such as range checks, size checks on both *isize* and *fsize*, etc., are taken, erroneous diagnostics can occur due to wrong data values in the super-block. *Isiz* (the contents of the first word in the super-block of the file system) is the block size of the *ilist* region. *Fsiz* (the contents of the second word in the super-block of the file system) is the block size of the entire file system.

Descriptions for each of the *error* mnemonics are given below:

- | | |
|-----|--|
| bad | The block number contains a value outside the allowable space on the file system. The block number cannot be less than (<i>isize</i> + 2) or greater than <i>fsize</i> . |
| dup | The block number is assigned to a previously-encountered file. This message represents only this particular reference to this block. The procedure for finding all such references is described in the next section on file patching. |
| din | The block is a member of a directory entry and contains an i-number outside the range defined by the <i>ilist</i> size of the file system. The size of the <i>ilist</i> is calculated by multiplying the value of <i>isize</i> by the number of inodes per disk block. There are currently 16 inodes in a 512 byte disk block. |
| blk | The block number was used as an argument (<i>-b</i> option) in <i>check</i> . |

The *block type* identifies the kind of block involved:

- | | |
|-------------|---|
| free | The block is not allocated. |
| data(small) | The block is a data block in a small file. A <i>small</i> file consists of eight or fewer disk blocks that directly contain the block addresses of data. |
| data(large) | The block is a data block in a large file. In a <i>large</i> file, each of the eight block addresses may reference an <i>indirect</i> block of up to 256 addresses of blocks that, in turn, constitute the file itself. |
| indirect | The block is an indirect block in a large file. |

Examples of these types of diagnostics messages are:

```
# check -b 53 /dev/rp0
/dev/rp0:
53 blk; inode= 0(free)
53 bad; inode= 0(free)
6 dups in free
3400 missing
spcl      159
files     254
large     88
direc     18
indir     88
used      3067
last      6681
free      92
```

```
# check /dev/rp54
/dev/rp54:
36638 dup; inode= 2205(data (small))
14563 dup; inode= 2208(data (small))
58318 din; inode= 2208(data (small))
58318 din; inode= 2208(data (small))
spcl      0
files     4970
large     742
direc     332
indir     767
used      44412
last      64666
free      19658
```

3.3 Inode Diagnostics

The third class of messages that will be discussed involves individual files, for which the actual number of directory references does not agree with the value contained in their link-count field in the super-block of the file system. The format of these messages is as follows:

`i#` error-indicator

The `#` represents the `i`-number of the inode. The error indicator consists of a byte whose octal value ranges from 0 to 377₈. The algorithm that `check` employs for computing the error indicator is explained below:

- a. 100₈ is added to the error indicator if the particular inode is allocated. The inode is disregarded if it is not allocated and its `mode` (permissions of the file) is zero.
- b. An additional 100₈ is added to the error indicator if the link-count field of the inode is non-zero; the value of the link-count field is also added to the error indicator.
- c. For each directory entry which references that particular inode, 1 is *subtracted* from the error indicator.

Only inodes whose error-indicators are neither 0 (unallocated) nor 200 (consistent file) are reported and flagged as requiring additional attention. The most common occurrences contain error-indicators of 100, 177, 201, 377. Example:

```
# check /dev/rp54
/dev/rp54:
    15  100
    23  201
    39  177
   213  377
spcl      0
files    4970
large    742
direc    332
indir    767
used     44412
last     64666
free     19658
```

Because of their frequent occurrence, a further explanation of the four diagnostics in the above example are given below.

3.3.1 Type-100 Error. In the case where the error-indicator is 100, the inode usually is allocated, but its link count is zero. This can be made evident by retracing through the steps that the *check* algorithm employs for computing the error-indicator value and by the example below, which illustrates the structure of such an inode in raw data format. *File System(V)* describes the complete structure of an inode and should be consulted for a more complete explanation. Example (see *fsdb(VIII)*):

```
# fsdb /dev/rp0
15i
i#: 15 md:a-----r----- ln: 0 uid: 0 gid: 0 s0: 0 s1: 0
a0: 1562 a1: 0 a2: 0 a3: 0 a4: 0 a5: 0 a6: 0 a7: 0
at: Thu May 26 13:31:06 1977 mt: Thu May 26 13:31:06 1977
```

where *md* is the mode (indicating allocation) and *ln* represents the link-count field.

This particular error is not catastrophic to the file system and does not necessitate immediate attention. The data blocks are still allocated, e.g., 1 block at 1562. The file is not referenced by any directory; hence, it is nameless. This error normally occurs as a result of a pipe being active when the system crashes.

3.3.2 Type-201 Error. Error-indicator 201 implies that an i-number does not appear in enough directories (201—missing in one directory, 202—missing in two directories, etc.). Such an occurrence is similar to the previous case in that its nature is not serious and does not warrant immediate attention. The data blocks associated with the inode will remain intact, but will not be available until the problem is remedied. Example:

```
# fsdb /dev/rp0
23i
i#: 23 md: a----rw-----ln: 1 uid: 11 gid:1 s0: 0 s1: 3050
a0: 660 a1: 671 a2: 724 a3: 824 a4: 34 a5: 1056 a6: 0 a7: 0
at: Thu May 12 16:05:10 1977 mt: Wed May 25 10:53:11 1977
```

The eight address blocks (*a0-a7*) contain the actual raw data for this particular example. A more detailed description of these eight address blocks can be found in *File System(V)*.

3.3.3 Type-377 Error. Error-indicator 377 (namely, -1 represented in two's complement format) indicates that an unallocated inode appears as an entry in a directory. This situation is *serious* and should be taken care of immediately. The major implication in this case is that the inode is not allocated and the link count is zero. This means that the system might reassign the inode to a new file and spread the infection to other parts of the file system. Example:

```
# fsdb /dev/rp0
213i
i#: 213 md: 0 ln: 0 uid: 11 gid: 1 s0: 0 s1: 0
a0: 0 a1: 0 a2: 0 a3: 0 a4: 0 a5: 0 a6: 0 a7:0
at: Thu May 12 16:05:10 1977 mt: Wed May 25 10:53:11 1977
```

where *md* indicates no allocation and *ln* represents the link count.

3.3.4 Type-177 Error. Error-indicator 177 indicates that the inode appears as an entry in one more directory than indicated by its link count. This problem, like the above, is serious and immediate attention should be given to correcting it before the infection spreads to other parts of the file system. For instance, removing an inode that has a 177 error indicator could turn it into one with a 377 error indicator. Example:

```
# fsdb /dev/rp0
39i
i#: 39 md: ad---rwxr-xr-x ln: 1 uid: 11 gid: 1 s0: 0 s1: 160
a0: 753 a1: 0 a2: 0 a3: 0 a4: 0 a5: 0 a6: 0 a7: 0
at: Thu May 12 16:05:10 1977 mt: Fri June 25 10:53:11 1977
```

Where *ln* represents the link-count field.

4. REPAIRING FILE SYSTEM DAMAGE

This section provides guidance and some methods for resolving the problems that lead to the various diagnostic messages discussed in Section 3 above. The solutions presented here are examples of actual occurrences.

4.1 Free-List Diagnostics

The first class of diagnostic messages are Free List Diagnostics. These types of messages are the most common and the least complex. They basically involve the portion of a file system known as the super-block. The super-block contains the sizes of both the *isize* and *fsize* and system tables of up to 100 free blocks and free i-numbers. Also, a pointer to a linked list of blocks, the free list, is contained within the super-block. The "root" (/) file system's super-block is considered as a special case, because the root file system is always mounted whenever PWB/UNIX is running. When a file system is mounted as a file system, a copy of its super-block is loaded into memory. Normally, before any actual repair work can be performed, the system should be set to the single-user environment. This will ensure no outside interference, as well as prevent spreading the infection to other file systems, because the root file system is normally the only file system mounted in single-user mode.

All repair work on a file system is made in its place of residence (normally a disk pack). Therefore, upon completion of repair work on the free list of the *root* file system, an *immediate reboot* of the system is required *without* a *sync(1)*. The reboot updates the in-memory super-block of the root file system with the newly-repaired copy. Omitting this step guarantees disastrous results.

4.1.1 Root File System. We show below a simple example involving the root file system and the necessary repair steps that are to be followed:

1.

```
# check /dev/rp0
/dev/rp0:
1 dups in free
1 missing
spel      159
files     270
large     97
direc     18
indir     98
used      3443
last      6681
free      3109
```

The above represents the present condition of the "root" file system after the system was rebooted and is in single-user mode. The error messages (dups in free, etc.) can now be repaired by invoking the command *icheck*, which searches through the specified file system and reconstructs the free list. The syntax for *icheck* is as follows:

```
# icheck -s[drive type] special-file
```

where:

s Indicates that the free list is to be built.

drive type Where the drive type corresponds to the numbers 3 through 6 (*RP03*, *RP04*, ... *RP06*) and, depending upon which type is specified, the free blocks are then built in that order.

special file File system device name (i.e., /dev/rp0, etc.).

The steps needed to repair the "root" file system in the above condition are as follows:

2.

```
# icheck -s4 /dev/rp0
/dev/rp0:
:
```

3.

```
# #00 = UNIX
system pwbb0215
Restricted rights:
Use, duplication ...

Available user memory = 6917 × 32 words
```

4.

```
# check /dev/rp0
/dev/rp0:
spel      159
files     270
large     97
direc     18
indir     98
used      3443
last      6681
free      3109
```

In step 2 above, the command *icheck* is invoked to reconstruct the free list. Step 3 shows that the system is then *halted* and then *rebooted*, because the "root" file system was involved as explained above. The remaining step, a recheck of the file system, is performed in order to make sure that the repair succeeded. The count of free blocks between steps 1 and 4 remained the same because the error message "1 dups in free" offsets the succeeding error message "1 missing" in step 1.

This procedure should be followed for all possible combinations of error messages that pertain to the "free list". The only exception to this procedure concerns step 3 above, which involves a reboot of the system if the root file system is specified.

4.1.2 Non-root File System. Below is an example that shows this procedure used on a non-root file system:

1.

```
# check /dev/rp44
/dev/rp44:
43 missing
spel      0
files     9207
large     1178
direc     673
indir     1185
used      58807
last      64999
free      5433
```

The above output from *check* represents the condition of the file system before any correctional steps are applied. The steps needed to repair a "non-root" file system in the above condition are as follows:

2.

```
# icheck -s4 /dev/rp44
/dev/rp44:
:
```

3.

```
# check /dev/rp44
/dev/rp44:
spel      0
files     9207
large     1178
direc     673
indir     1185
used      58807
last      64999
free      5476
```

The 43 blocks that were shown as "missing" in step 1 are now part of the reconstructed free list.

4.2 Block Diagnostics

The next class of diagnostic messages is Block Diagnostics. Section 3 described the two distinct types that make up this category. Their respective formats are as follows:

1. b# error; inode = i# (block type)
2. i# error-indicator

This type requires checking each disk block in a file system that is associated with a particular inode.

4.2.1 "DUP" Diagnostic. A typical example is the error message *dup* followed by an i-number and the type of disk block. We now show a procedure used to resolve this kind of problem:

1.

```
# check /dev/rp2
/dev/rp2:
25511 dup; inode = 2201 data (small)
spcl      0
files    3987
large    858
direc    338
indir    883
used     39647
last     52660
free     12952
```

2. The initial check of the file system shows a *dup* error message; this message indicates that the allocated block number (25511) contains more than one *reference point*. This reference point can refer to either a previous file or possibly the same file. The first step is to find all the reference points (i-numbers of files), that are referencing this allocated disk block. This is done by:

```
# check -b 25511 /dev/rp2
```

This will give a complete listing of all the i-numbers that *reference* the disk block 25511.

```
/dev/rp2:
25511 blk; i= 1105 data(small)
25511 blk; i= 2201 data(small)
25511 dup; i= 2201 data(small)
spcl      0
files    3987
large    858
direc    338
indir    883
used     39647
last     52660
free     12952
```

3. The above output shows two i-numbers, 1105 and 2201, that reference the disk block 25511. The names associated with these i-numbers can be determined by invoking the command *ncheck*(VIII). *Ncheck* lists the given i-numbers with their respective names:

```
# ncheck -i 1105 2201 /dev/rp2
/dev/rp2:
1105     /hasp/info/logx484
2201     /hasp/testjob
#
```

4. One now has to determine which of the two i-numbers is the rightful owner of the disk block 25511. Sometimes this is impossible; the only possible solution is to remove *all* the files named in step 3. In this particular example, uncertainty is not involved; the data contents of disk block 25511 are displayed and give clues in identifying the correct owner, as follows:

```
# fsdb /dev/rp2
25511b,64c
61647000: / / t e s t j o b   j o b   ( 6
61647020: 2 7 0 , r 8 4 0 ) , r j e ,   c l
61647040: a s s = x\n / / *           u s r
61647060: = ( h a s p , j o b )\n / / i e
```

The first line above invokes *fsdb*(VIII) with the names of the file system as its argument. *Fsdb* provides an efficient method for scanning an inode or a disk block. The second line is supplied by the user as input to *fsdb*; the "25511" is the (decimal) number of the block that will be inspected. The "b" expands the block number to a full block address, (i.e., 61647000). The "64" indicates that 64 characters are to be printed. The letter "c" causes the 64 characters to be printed in *character mode*. The output resembles the contents of the file /usr/hasp/testjob.

5. Once the correct owner has been established, the file can then be removed, but only after the file system /dev/rp2 is *mounted*. Steps 1 through 4 were accomplished while the file system /dev/rp2 was *unmounted* or *inactive*:

```
# mount /dev/rp2 /mnt
WARNING!! - mounting: <c2> as </mnt>
# rm -f /mnt/hasp/info/logx484
```

At this point, the file system "/c2" can then be *unmounted* (first line below), its *free* list reconstructed (next two lines), and checked again:

```
# umount /dev/rp2
# icheck -s4 /dev/rp2
/dev/rp2:
:
# check /dev/rp2
/dev/rp2:
spel          0
files         3986
large         858
direc         338
indir         883
used          39646
last          52660
free          12952
```

The process of identifying the principal owner (i-number) of the disk block sometimes fails. When this occurs, one possible solution is to remove *all* files that reference that disk block (25511). The procedure in this case remains the same as before, with the exception that step 4 and the last line of step 5 are replaced with the following:

```
# rm -f /mnt/hasp/info/logx484
# rm -f /mnt/hasp/testjob
```

All files that are removed should be restored from the latest backup medium for that file system. The owner of each such file must then be notified of the problem and be given the full path name of each restored file.

An alternative procedure would be to first recopy all the files involved to *temporary* files, and then remove the original files. The owners should be notified that their files have been renamed and that they possibly contain contaminated information.

4.2.2 An Example of a More Complex "DUP" Diagnostic. The above case represents a case that does not involve massive damage. However, when massive damage is present, the repair process is substantially more complex, as shown in the following example:

```
# check /dev/rp34
/dev/rp34:
16705 dup; inode = 12960 (data (large))
8270 dup; inode = 12960 (data (large))
8224 dup; inode = 12960 (data (large))
13112 dup; inode = 12960 (data (large))
13873 dup; inode = 12960 (data (large))
13360 dup; inode = 12960 (data (large))
8224 dup; inode = 12960 (data (large))
12320 dup; inode = 12960 (data (large))
21332 dup; inode = 13052 (data (large))
8242 dup; inode = 13991 (data (small))
12342 dup; inode = 15175 (data (small))
18516 dup; inode = 15220 (data (small))
16967 dup; inode = 16164 (data (large))
21326 dup; inode = 16948 (data (small))
19785 dup; inode = 17305 (data (large))
20037 dup; inode = 17305 (data (large))
1 dups in free
14 missing
spel          0
files        15428
large         492
direc         701
indir         525
used          52193
last          64989
free          11485
```

The inode 12960 claims to be the owner of several disk blocks. This inode is the most likely candidate that should be removed in order to restore consistency to the file system. The following steps will repair the file system damage and restore the file system to a working condition:

```
# ncheck -i 12960 /dev/rp34
/dev/rp34:
12960 /9452/sita/pscsxs/aaaacsh1a/AAN---831604--
# mount /dev/rp34 /mnt
WARNING!! - mounting: <b4> as </mnt>
# rm -f /mnt/9452/sita/pscsxs/aaaacsh1a/AAN---831604--
# umount /dev/rp34
# icheck -s4 /dev/rp34
/dev/rp34:
:
:
# check /dev/rp34
/dev/rp34:
spel          0
files        15427
large         491
direc         701
indir         524
used          51936
last          64989
free          11500
```

The output of the last *check* command above reveals that the damaged file system, /dev/rp34, was repaired and is now consistent. However, the procedure cannot guarantee that the data blocks associated with the other affected files contain valid information. The procedure can be expanded to include the necessary steps to identify the path names of all the affected files in the above list. The owners can

then be notified that their files may contain invalid data. The additional steps replace the *ncheck* command in the above procedure.

Each block number reported as a *dup* should be checked. This will then identify all the *i*-numbers referencing a duplicate block:

```
# check -bbb ...b 16705 8270 8224 ... 20037 /dev/rp34
```

Where *n* occurrences of *b* (-bbb...b) cause the next *n* arguments to be interpreted as block numbers. Each unique *i*-number reported in the previous step can then be converted into a path name.

```
# ncheck -i [i-numbers] /dev/rp34
```

4.2.3 "BAD" Diagnostic. The message *bad* implies that the block number contains a value outside the *allowable space* on the file system. *Allowable space* is defined as the number of disk blocks logically addressable, from *isize* + 2 to the limit specified by the super-block entry *fsize*.

These messages usually occur for a small file that was in the process of being transformed into a large file at the time of the crash. The mode field in the inode control structure indicates a large file, but the address block references a data block address instead of an indirect block address.

Because the bad block and the file are normally unrelated to one another, the contaminated file must be removed. The example below will be used as a model for developing a procedure for this kind of problem:

```
# check /dev/rp24
/dev/rp24:
49802 bad; inode = 1191 (data(large))
49810 bad; inode = 1191 (data(large))
128 bad; inode = 1191 (data(large))
32768 dup; inode = 1191 (data(large))
32768 dup; inode = 1191 (data(large))
40960 bad; inode = 1191 (data(large))
1 bad; inode = 1191 (data(large))
436 bad; inode = 1191 (data(large))
4 dups in free
73 missing
spel      0
files    4378
large    854
direc    299
mdir     859
used    31218
last    65475
free     4118
```

Once the initial check of the file system has been performed, as above, the following steps must be performed:

- 1 Unallocate the *i*-number 1191. This will unallocate the bad blocks and leave the *i*-number as a non-existent file with its directory entry and name field intact. This is done by using *clrm*(VIII), (rather than *rm*(I), because the system checks for bad block numbers when *rm* is used, and might not remove the file):

```
# clrm /dev/rp24 1191
```

- 2 Check the file system to insure that the bad messages were reduced into a simpler diagnostic:

```
# check /dev/rp24
/dev/rp24:
4 dups in free
14 missing
  1191 377
spel      0
files    4377
large    853
direc    299
indir    858
used    31215
last    35989
free     4118
```

3. Convert the i-number 1191 from step 2 into a path name:

```
# ncheck -i 1191 /dev/rp24
/dev/rp24:
1191 /adm/core
```

4. Mount the file system:

```
# mount /dev/rp24 /mnt
WARNING!! - mounting: <a1 > as </mnt >
```

5. Remove the name of the non-existent file:

```
# rm -f /mnt/adm/core
```

6. Unmount the file system:

```
# umount /dev/rp24
```

7. Reconstruct the free list to eliminate the "dups in free" that result from step 5.

```
# icheck -s3 /dev/rp24
```

8. Check the file system again:

```
# check /dev/rp24
/dev/rp24:
spel      0
files    4377
large    853
direc    299
indir    858
used    31215
last    35989
free     4129
```

9. The contaminated file that was removed in step 5 should be restored from the latest back-up medium for that file system and its owner notified.

4.2.4 "DIN" Diagnostic. The message *dn* means that a block number is a member of a directory entry containing an i-number that resides outside the range of the *list* size of the file system; *size* (the contents of the first word in the super-block of the file system) is the block size of the *list* region.

These messages normally occur when a directory contains an address block that references text data instead of directory entries. A *directory entry* consists of an i-number and a 14-character name. The example below will be used as a model for developing a procedure for this kind of problem.

```
# check /dev/rp24
/dev/rp24:
7902 din; inode = 15 (data(small))
1 dups in free
1 missing
  1 201
  2 201
  3 201
  5 201
  7 201
  9 201
 13 201
 14 201
 15 201
 33 201
 168 201
 169 201
 172 201
 340 201
 488 201
 808 201
2975 201
3976 201
3977 201
4099 201
4335 201
4364 201
4547 201
5002 201
8202 377
8224 342
spel      0
files     4769
large     870
direc     251
indir     879
used     33545
last     35989
free      1798
```

The first step is to position *fsdb* at the beginning of the inode structure for i-number 15 in file system */dev/rp24*.

```
1. # fsdb /dev/rp24
15i
i#: 15 md:ad---rwxr-xr-x ln: 3 uid: 2 gid: 0 s0: 0 s1: 560
a0: 7902 a1: 3314 a2: 0 a3: 0 a4: 0 a5: 0 a6: 0 a7: 0
at: Fri May 27 17:30:02 1977 mt: Fri May 27 07:01:02 1977
```

Where 15i is the offset into the file system. The first address field (a0) contains the same block address as was reported by the initial *check* above. The block 7902 and its data contents will be further investigated.

```
2. 7902b.p32d
d0: 8224      Field.
d1: 8224
d2: 8224
d3: 8224
```

```
d4: 8224
d5: 8224      IF
d6: 8224
d7: 8224
d8: 8224
d9: 8224
d10: 8224     Only
d11: 8224
d12: 8224
d13: 8224
d14: 8224
d15: 28448    NE/n
d16: 8224
d17: 8224
d18: 8224
d19: 8224     PM
d20: 8202
d21: 8224
d22: 8224
d23: 8224
d24: 8224     IS
d25: 8224
d26: 8224
d27: 8224
d28: 8224
d29: 8224     DESIRED
d30: 8224
d31: 8224
q
```

In the above, the "7902b.p32d" displays the contents of block 7902 as directory entries; q terminates this command. The block 7902 is obviously a data block and not a block of directory entries. The number 8224 in each instance is *interpreted* as an i-number and accounts for the diagnostic message 342 (last diagnostic in the output of the *check* command above). The *dm* diagnostic occurred at the fifteenth entry (d15:, in step 2 above) because the number, 28448, exceeds the boundary limitations of the *ilist size* (the *ilist size* for file system, /dev/rp24, was defined as 655 blocks, with each block containing a maximum of 16 inodes for a maximum of 10480 i-numbers). This shows that the directory block is damaged. Because directories change infrequently, as a rule, there is, in all probability, a good copy on the most recent backup pack.

The next step is to physically mount (on another drive) the latest backup copy for the file system /dev/rp24 and then to replace the contaminated block address (7902) on the original file system with the corresponding inode's block address from the backup copy.

3. # fsdb (backup file system)

```
15i
i#: 15 md: ad---rwxr-xr-x ln: 3 uid: 2 gid: 0 s0: 0 s1: 560
a0: 1344 a1: 3314 a2: 0 a3: 0 a4: 0 a5: 0 a6: 0 a7: 0
at: Fri May 27 07: 01: 01 1977 mt: Fri May 27 07: 01: 01 1977
```

Note that we used the same offset (15i) into the backup file system as we did above for the damaged file system.

4. The contents of the first address block (1344) in Step 3 are displayed below:

```
1344b.p32d
d0: 15
d1: 1 ..
d2: 14 dtablezero
d3: 13 SA
d4: 5002 dtable
d5: 808 lock
d6: 0 dtbsv
d7: 9 dailyacct
d8: 4099 diskusage
d9: 7 wtmp
:
:
d31: 0 x
q
```

In the above, the "1344b.p32d" displays the contents of block 1344 as directory entries; q terminates this command. Note the difference between the above output and the output in step 2; also the same i-numbers that were reported as 201 type errors in step 2 will have valid directory entries once the block transformation is made.

The next step is to determine if the block number (1344) on the original file system is unallocated. This step is necessary in order to prevent duplicate block numbers in files. The following command line will perform this function:

```
# check -b 1344 /dev/rp24 (original file system)
```

If the response from the above command is

```
1344 blk; inode = 0 (free)
```

then block number 1344 can be used. Otherwise, the block is already allocated and an alternate block number must be used from the free list.

For the case where the same block number is to be reused, there is a high probability that the contents within the block have not been altered and thus eliminates the need for recopying the block before we use it. The block 1344 on the original file system should be inspected; its contents should resemble the data contents in Step 4. If differences between the contents of the two blocks are present, it is recommended that the block be re-copied with *bcopy*(VIII) before using the block. However, the contaminated block address (7902) on the original file system must be changed to the new block address (e.g.,1344):

```
5. # fsdb /dev/rp24 (original file system)
15i.a0 = 1344
q
```

The steps below will reconstruct the free list and re-check the file system for accuracy:

```
6. # icheck -s3 /dev/rp24
```

```
7. # check /dev/rp24
```

```
/dev/rp24:
spcl      90
files    4769
large    870
direc    251
indir    879
used    33545
last    35989
free     1798
```


4.3 Inode Diagnostics

The last class of error diagnostics that will be discussed here has the following format:

```
i# error-indicator
```

Section 3 defined the framework for this category of diagnostics and identified the types that occur most frequently. All instances will now be examined separately and will illustrate methods that can be employed to repair them. Although each of these types of errors is different in definition and complexity, they do share some similarity with respect to their repair processes.

4.3.1 Type-100 Error. We first discuss the "type-100" errors ("type-201" errors are discussed in Section 4.3.4). The example below shows both 100 and 201 type errors.

```
# check /dev/rp0
/dev/rp0:
  426 100
  427 201
  429 100
  431 201
  435 100
  438 100
spcl   174
files  268
large   92
direc   29
used   2904
last   6681
free   3661
```

The above example pertains to the "root" file system. This fact is important, because the root file system is a special case: it requires an immediate reboot of the system after any alterations are made in its super-block.

The "type-100" error usually implies that the inode is allocated, but that its link-count is zero. The process for repairing "type-100" errors is to make the inode unallocated. The command *clrm*(VIII) (clear inode mode) performs this function. The procedure is as follows:

1. *clrm* file-system [i-numbers ...]

e.g.:

```
# clrm /dev/rp0 426 429 435 438
```

2. # *icheck* -s4 /dev/rp0

```
/dev/rp0:
```

```
⋮
```

```
# #00= unix
system pwbd0127
Restricted rights:
```

```
⋮
```

```
Available user memory = 6910 × 32 words
```

```
3. # check /dev/rp0
   /dev/rp0:
     427  201
     431  201
   spl      174
   files    264
   large    92
   direc    29
   used     2904
   last     6681
   free     3670
```

The first command above reconstructs the root free list; because the root file system is involved, the system is then rebooted in single-user mode, and re-checked to ensure that all "type-100" errors have been taken care of. This procedure can be applied to any *inactive* file system. The rebooting after the *check* is needed when the "root" file system is involved.

4.3.2 *Type-377 Error.* The next type in this class of diagnostic messages consists of "type-377" errors (-1 represented in two's complement form). This particular error type indicates that an unallocated inode is represented as an entry in a directory as shown in the example below:

```
# check /dev/rp14
/dev/rp14:
37 missing
  1992 377
spl      0
files    11364
large    1069
direc    817
indir    1074
used     47872
last     64998
free     15864
```

Because this type of error is considered highly communicable in spreading the infection to other files, it is crucial that this be attended to immediately. This error can also cause diagnostics with error indicator values of 376, 375, etc. Instances where the value of the error indicator is other than 377 (i.e., 376, 375, etc.), are most likely to be *directories* or *multiply-linked files*. When these conditions arise, the repair procedures become more complex and demand a greater amount of knowledge and experience. These errors are also very serious and warrant immediate attention.

The repair process for a multiply-linked file is identical to the procedure for repairing an ordinary file that contains an error indicator value equal to 377_g. The differences are that several file names replace the single name referenced in both step 1 and step 3 of the procedure below and that the error indicator should equal values of 376, 375, etc.

When a directory is involved, the error condition becomes compounded, because the directory itself and its contents (*files* and *subdirectories*) are also affected. The first step is to determine the complete path name of the directory. Once the path name has been determined, the entire directory can then be removed and later restored from the latest available backup medium.

The process of converting an i-number into a path name is accomplished by the command *ncheck*. This command will terminate prematurely when a directory that is needed to complete the construction of the path name is non-existent. The path name must then be constructed manually, which is difficult and error-prone.

The procedure for repairing an ordinary file that contains an error indicator of 377 is as follows.

1. The first step is to convert the i-number (1992) into a character name:

```
# ncheck -i 1992 /dev/rp14
/dev/rp14:
1992 /f1/payload/r43117
```

2. Once the path name is determined, the file system must now be mounted as an active file system:

```
# mount /dev/rp14 /mnt
WARNING!! - mounting: <f3> as </mnt>
```

3. The name of the non-existent file can now be removed:

```
# rm -f /mnt/f1/payload/r43117
```

4. The file system is then unmounted:

```
# umount /dev/rp14
```

5. The file system is checked again for accuracy:

```
# check /dev/rp14
/dev/rp14:
spcl          0
files        11364
large         1069
direc         817
indir         1074
used         47872
last         64998
free         15901
```

4.3.3 *Type-177 Error.* The error indicator value of 177, 176, 175, etc., indicate that an inode is referenced as an entry by more directories than its actual link count indicates. This problem, similar to the type above, is serious, and should be corrected immediately. An example of this error type is shown below:

```
# check /dev/rp32
/dev/rp32:
612 177
spcl      76
files     288
large     101
direc     18
indir     101
used     3236
last     4397
free     1098
```

The correction procedure requires an adjustment in the inode's link-count field. The exact amount is the difference between the constant, 200_8 and the inode's error-indicator. In this particular case, the difference between the two values is 1 ($200_8 - 177_8$). The value 200_8 represents a file having no abnormal characteristics. The error-indicator 177 represents the present status of the file. The most logical action is to *increase* the link-count field by a value equal to the difference between these two numbers. The procedure for altering the link-count field is as follows:

```
# fsdb /dev/rp32
612i
i#: 612 md: ad---rwxr-xr-x ln: 1 uid: 0 gid: 0 s0: 0 s1: 32
a0: 891 a1: 0 a2: 0 a3: 0 a4: 0 a5: 0 a6: 0 a7: 0
at: Wed Mar 16 13:36:42 1977 mt: Sat Feb 19 17:39:55 1977
ln= 2
q
# check /dev/rp32
/dev/rp32:
spcl      76
files     288
large     101
direc     18
indir     101
used      3236
last      4397
free      1098
```

The above procedure first seeks to the starting address of the control block for i-number 612, then sets the link-count field to 2 (1 more than it was) and finally re-checks the file system to insure the procedure worked correctly.

4.3.4 Type-201 Error. The last type of error diagnostic consists of the error indicator 201. This type of error implies that a directory entry is "missing"; it is *not* serious and does *not* require immediate attention. However, the data blocks associated with the inode will remain unavailable until the problem is solved. The "missing" directory entry can be either an *ordinary file* or a *directory*. Therefore, the first step involves identifying the i-number as either a directory or an ordinary file. This identification process is accomplished by examining the mode field in the inode structure for the particular i-number. If the inode is identified as a directory, its repair procedure is deferred until the end of this section. Otherwise, the inode is an ordinary file and its data blocks can be salvaged by copying the inode's i-number in an empty entry in some directory. The directory which is involved in this process is called "201dir"; it resides as a top-level directory under the root directory of each file system on all of the PWB/UNIX Systems. An example of this error type is shown below:

```
# check /dev/rp0
/dev/rp0:
 427 201
 431 201
spcl   174
files  264
large   92
direc  29
mdir   92
used   2903
last   6681
free   3650
```

The repair process for this error type is different from all the previous procedures because a single-user environment is not required. The repair procedure is as follows:

1. The initial step is to identify the i-number of the directory "201dir":

```
# ls -id /201dir
92 /201dir
# chdir /201dir
```

2. We are now able to locate ourselves at that position in the */usr* of the root file system:

```

# fsdb /dev/rp0
92i
#i: 92 md: ad---rwxr-xr-x ln: 2 uid: 0 gid: 0 s0: 0 s1: 112
      *
a0: 2480 a1: 0 a2: 0 a3: 0 a4: 0 a5: 0 a6: 0 a7: 0
at: Wed Mar 30 11:57:12 1977 mt: Mon Feb 28 07:40:09 1977

```

3. The next step is to position ourselves in the file system at address `addr[0]` in the `201dir`'s inode or control block structure. The value of `addr[0]` is indicated by the "*" in the above inode structure.

```

a0b.p4d
d0: 92
d1: 1 ..
d2: 0 446
d3: 0 x1

```

The "a0b" is a mnemonic notation for `addr[0]`. The "p4d" prints 4 entries in the "201dir" directory. The decimal number 92 is the i-number of the "201dir" directory. The "." represents the current directory (the result of step 1). The ".." represents the parent directory of "201dir" its i-number is 1, which represents the "root" directory or "/". The remaining two entries (d2 and d3) are currently null. Their names are *446* and *x1*, respectively. Notice their i-numbers are both zero and currently the files themselves are non-existent. These two empty entries, *446* and *x1*, will soon contain the i-numbers of the two 201 errors.

```

a0b.d2= 427
d2: 427

```

4. We assign the i-number 427 to the null entry 446, which now becomes the name of that *orphan*. Its full path name is `/201dir/446`. We repeat the procedure for the *orphan* i-number 431:

```

a0b.d3= 431
d3: 431

```

5. We now terminate the command invoked in step 2 and identify the respective owners and sizes of the two files `/201dir/446` and `/201dir/x1`, and restore them to their owners.

```

# ls -l 446 x1
-rw-----1 deem 3072 Feb 9 08:38 446
-rw-----1 atrs 3072 Feb 9 08:37 x1

```

The name field of a directory entry, other than "." and "..", can be changed to represent any character string. The string is limited to a maximum of 14 characters. For instance, the present name of directory entry d2 in `/201dir` is *446*; the following step will change the name from *446* to *orphan446*:

```

a0b.d2.nm = orphan446

```

The procedure in steps 1 through 5 above assumes the presence of empty (null) entries in the "201dir" directory. If these empty entries are not present, the user must then make them available before steps 4 and 5 can be executed. Step 3 notifies the user of the existence of such null entries by printing their names. If they do not appear, the following steps will create these missing null entries:

```

q
# cp /dev/null x
# cp x x1
# cp x1 x2
# rm x x1 x2

```

The above terminates the `fsdb` command of step 2 and creates three empty files in `201dir`, (*x*, *x1*, *x2*), which are then immediately removed consequently producing the null entries in `201dir`. The repair procedure can now be restarted at step 2. This procedure is not complex but is very susceptible to human mistakes; therefore, it is necessary that each step be thoroughly understood.

The affected i-number of a "type-201" error, as well as the previous error types (177, 100, 377, etc), can apply to a *directory*, a *special-file*, or *no file at all*, in addition to an *ordinary file*. The various types of

PWB/UNIX files and their description are given in [1]. The first word in the i-number's inode structure identifies the file type and indicates inode allocation and file permissions:

file type:	
directory	d
character device	c
block device	b
large file	l
ordinary file	-
inode allocation	a
file permissions:	
owner	rx
group	rx
others	rx

e.g.,

```
fsdb /dev/rp0
15i
i#: 15 md: ad---rwxr-xr-x ln: 2 uid: 11 gid: 1 s0: 0 s1: 3050
a0: 660 a1: 671 a2: 724 a3: 824 a4: 34 a5: 1056 a6: 0 a7: 0
at: Thu May 12 16:05:10 1977 mt: Wed May 25 10:53:11 1977
q
```


When a "type-201" error is identified as a directory, the repair process is different than for an ordinary file. The procedure is similar to that for resolving a "177" error diagnostic. The directory's link count must be *decreased* by a value equal to the difference between the error-indicator value (201_g) and the constant 200_g, namely by 1.


5. ACKNOWLEDGEMENTS

I wish to thank Rich Graveman and Larry Wehr for suggestions on clarifying and simplifying some of the more complex procedures. I would also like to express my gratitude to T. A. Dolotta for reviewing several drafts of this document and to Margarita Rivera for her help in its preparation.

6. REFERENCES

- [1] D. M. Ritchie and K. Thompson, The UNIX Time-Sharing System, *Comm. ACM* **17**(7):365-75 (July 1974).
- [2] T. A. Dolotta, R. C. Haight, and E. M. Piskorik, eds., *PWB/UNIX User's Manual—Edition 1.0*, Bell Laboratories (May 1977)







The PWB/UNIX* document entitled:
PWB/UNIX RJE Administrator's Guide
is not yet available.

* UNIX is a Trademark/Service Mark of the Bell System.







The Unix I/O System

Dennis M. Ritchie
Bell Telephone Laboratories

This paper gives an overview of the workings of the Unix I/O system. It was written with an eye toward providing guidance to writers of device driver routines, and is oriented more toward describing the environment and nature of device drivers than the implementation of that part of the file system which deals with ordinary files.

It is assumed that the reader has a good knowledge of the overall structure of the file system as discussed in the paper "The Unix Time-sharing System." Moreover the present document is intended to be used in conjunction with a copy of the system code, since it is basically an exegesis on that code.

Device Classes

There are two classes of device: *block* and *character*. The block interface is suitable for devices like disks, tapes, and DECTape which do, or can, work in 512-byte blocks and can be used in direct-access fashion. Ordinary magtape just barely fits in this category. Block devices can at least potentially contain a mounted file system. The interface to block devices is very highly structured; the drivers for these devices share a great many routines as well as a pool of buffers.

Character-type devices have a much more straightforward interface, although more work must be done by the driver itself.

Devices of both types are named by a *major* and a *minor* device number. These numbers are generally stored as a word with the minor device number as the low byte and the major device number as the high byte. The major device number selects which driver will deal with the device; the minor device number is not used by the rest of the system but is passed to the driver at appropriate times. Typically the minor number selects a subdevice attached to a given controller, or one of several similar hardware interfaces.

The major device numbers for block and character devices are used as indices in separate tables; they both start at 0 and therefore overlap.

Overview of I/O

The purpose of the *open* and *creat* system calls is to set up entries in three separate system tables. The first of these is the *u_file* table, which is stored in the system's per-process data area *u*. This table is indexed by the file descriptor returned by the *open* or *creat*, and is accessed during a *read*, *write*, or other operation on the open file. The only contents of the entry is a pointer to the corresponding entry of the *file* table, which is a per-system data base. There is one entry in the *file* table for each instance of *open* or *creat*. This table is per-system because the same instance of an open file must be shared among the several processes which can result from *forks* after the file is opened. A *file* table entry contains flags which indicate whether the file was open for reading or writing or is a pipe; a count which is used to decide when all processes using the entry have terminated or closed the file and therefore when the entry can be released. There is also a 32-bit file offset which is used to indicate where in the file the next read or write will take place. Finally, there is a pointer to the entry for the file in the *inode* table, which contains a copy of the file's i-node. Notice that an entry in the *file* table corresponds precisely to an instance of *open* or *creat*; if the same file is opened several times, it will have several entries in this table. However, there is at most one entry in the *inode* table for a given file. Also, a file may enter the *inode* table not only because it is open, but also because it is the current directory of some process or because it is a special file containing a currently-mounted file system.

An entry in the *inode* table differs somewhat from the corresponding i-node as stored on the disk; the modified and accessed times are not stored, and the entry is augmented by a flag word containing information about the entry, a count used to determine when it may be allowed to disappear, and the device and i-number whence the entry came.

During the processing of an *open* or *creat* call for a special file, the system always calls the device's *open* routine to allow for any special processing required (rewinding a tape, turning on the data-terminal-ready lead of a modem, etc.). However, the *close* routine is called only when the last process closes a file, that is, when the i-node table entry is being deallocated. Thus it is not feasible for a device to maintain, or depend on, a count of its users, although it is quite possible to implement an exclusive-use device which cannot be reopened until it has been closed.

When a *read* or *write* takes place, the user's arguments and the *file* table entry are used to set up the variables *u.u_base*, *u.u_count*, and *u.u_offset* which respectively contain the (user) address of the I/O target area, the byte-count for the transfer, and the current location in the file. If the file referred to is a character-type special file, the appropriate read or write routine is called; it is responsible for transferring data and updating the count and current location appropriately as discussed below. Otherwise, the current location is used to calculate a logical block number in the file. If the file is an ordinary file the logical block number must be mapped (possibly using an indirect block) to a physical block number; a block-type special file need not be mapped. In any event, the resulting physical block number is used, as discussed below, to read or write the appropriate device.

Character device drivers

The *cdevsw* table specifies the interface routines present for character devices. Each device provides five routines: *open*, *close*, *read*, *write*, and *special-function*. Any of these may be missing. If a call on the routine should be ignored, (e.g. *open* on non-exclusive devices which require no setup) the *cdevsw* entry can be given as *nulldev*; if it should be considered an error, (e.g. *write* on read-only devices) *nodev* is used.

The *open* routine is called each time the file is opened with the full device number as argument. The second argument is a flag which is non-zero only if the device is to be written upon.

The *close* routine is called only when the file is closed for the last time, that is when the very last process in which the file is open closes it. This means it is not possible for the driver to maintain its own count of its users. The first argument is the device number; the second is a flag which is non-zero if the file was open for writing in the process which performs the final *close*.

When *write* is called, it is supplied the device as argument. The per-user variable *u.u_count* has been set to the number of characters indicated by the user; for character devices, this number may be 0 initially. *u.u_base* is the address supplied by the user from which to start taking characters. The system may call the routine internally, so the flag *u.u_segflg* is supplied which indicates, if *on*, that *u.u_base* refers to the system address space instead of the user's.

The *write* routine should copy up to *u.u_count* characters from the user's buffer to the device, decrementing *u.u_count* for each character passed. For most drivers, which work one character at a time, the routine

```
cpass( )
```

is used to pick up characters from the user's buffer. Successive calls on it return the characters to be written until *u.u_count* goes to 0 or an error occurs, when it returns -1. *Cpass* takes care of interrogating *u.u_segflg* and updating *u.u_count*.

Write routines which want to transfer a probably large number of characters into an internal buffer may also use the routine

```
iomove(buffer, offset, count, flag)
```

which is faster when many characters must be moved. *Iomove* transfers up to *count* characters into the *buffer* starting *offset* bytes from the start of the buffer; *flag* should be *B_WRITE* (which is 0) in the write case. Caution: the caller is responsible for making sure the count is not too large and is non-zero. As an efficiency note, *iomove* is much slower if any of *buffer+offset*, *count* or *u.u_base* is odd.

The device's *read* routine is called under conditions similar to *write*, except that *u.u_count* is guaranteed to

be non-zero. To return characters to the user, the routine

```
passc(c)
```

is available; it takes care of housekeeping like *cpass* and returns -1 as the last character specified by *u.u_count* is returned to the user; before that time, 0 is returned. *lomove* is also usable as with *write*; the flag should be *B_READ* but the same cautions apply.

The "special-functions" routine is invoked by the *stty* and *gtty* system calls as follows:

```
sgtty(dev, v)
```

where *dev* is the device number and *v* is a vector. In the *gtty* case, the device is supposed to place up to 3 words of status information into the vector; this will be returned to the caller. In the *stty* case, *v* is 0 ; the device should take up to 3 words of control information from the array *u.u_arg[0...2]*.

Finally, each device should have appropriate interrupt-time routines. The interrupt-catching mechanism makes the low-order four bits of the "new PS" word in the trap vector for the interrupt available to the interrupt handler. This is conventionally used by drivers which deal with multiple similar devices to encode the minor device number.

A number of subroutines are available which are useful to character device drivers. Most of these handlers, for example, need a place to buffer characters in the internal interface between their "top half" (read/write) and "bottom half" (interrupt) routines. For relatively low data-rate devices, the best mechanism is the character queue maintained by the routines *getc* and *putc*. A queue header has the structure

```
struct {
    int    c_cc;    /* character count */
    char  *c_cf;    /* first character */
    char  *c_cl;    /* last character */
} queue;
```

A character is placed on the end of a queue by

```
putc(c, &queue)
```

where *c* is the character and *queue* is the queue header. The routine returns -1 if there is no space to put the character, 0 otherwise. The first character on the queue may be retrieved by

```
getc(&queue)
```

which returns either the (non-negative) character or -1 if the queue is empty.

Notice that the space for characters in queues is shared among all devices in the system and in the standard system there are only some 600 character slots available. Thus device handlers, especially write routines, must take care to avoid gobbling up excessive numbers of characters.

The other major help available to device handlers is the sleep-wakeup mechanism. The call

```
sleep(event, priority)
```

causes the process to wait (allowing other processes to run) until the *event* occurs; at that time, the process is marked ready-to-run and the call will return when there is no process with higher *priority*.

The call

```
wakeup(event)
```

indicates that the *event* has happened, that is, causes processes sleeping on the event to be awakened. The *event* is an arbitrary quantity agreed upon by the sleeper and the waker-up. By convention, it is the address of some data area used by the driver, which guarantees that events are unique.

Processes sleeping on an event should not assume that the event has really happened; they should check that the conditions which caused them to sleep no longer hold.

Priorities can range from 127 to -127 ; a higher numerical value indicates a less-favored scheduling situation. A process sleeping at negative priority cannot be terminated for any reason, although it is conceivable that it may be swapped out. Thus it is a bad idea to sleep with negative priority on an event which

might never occur. On the other hand, calls to *sleep* with non-negative priority may never return if the process is terminated by some signal in the meantime. Incidentally, it is a gross error to call *sleep* in a routine called at interrupt time, since the process which is running is almost certainly not the process which should go to sleep. Likewise, none of the variables in the user area "u." should be touched, let alone changed, by an interrupt routine.

If a device driver wishes to wait for some event for which it is inconvenient or impossible to supply a *wakeup*, (for example, a device going on-line, which does not generally cause an interrupt), the call

sleep(&lbolt, priority)

may be given. *Lbolt* is an external cell whose address is awakened once every 4 seconds by the clock interrupt routine.

The routines

spl4(), spl5(), spl6(), spl7()

are available to set the processor priority level as indicated to avoid inconvenient interrupts from the device.

If a device needs to know about real-time intervals, then

timeout(func, arg, interval)

will be useful. This routine arranges that after *interval* sixtieths of a second, the *func* will be called with *arg* as argument, in the style

*(*func)(arg)*

Timeouts are used, for example, to provide real-time delays after function characters like new-line and tab in typewriter output, and to terminate an attempt to read the 201 Dataphone *dp* if there is no response within a specified number of seconds. Notice that the number of sixtieths of a second is limited to 32767, since it must appear to be positive, and that only a bounded number of timeouts can be going on at once. Also, the specified *func* is called at clock-interrupt time, so it should conform to the requirements of interrupt routines in general.

An example

The driver for the paper-tape reader/punch is worth examining as a fairly simple example of many of the techniques used in writing character device handlers. The *pc11* structure contains a state (used for the reader), an input queue, and an output queue. A structure, rather than three individual variables, was used to cut down on the number of external symbols which might be confused with symbols in other routines.

When the file is opened for reading, the *open* routine checks to see if its state is not *CLOSED*; if so an error is returned since it is considered a bad idea to let several people read one tape simultaneously. The state is set to *WAITING*, the interrupt is enabled, and a character is requested. The reason for this gambit is that there is no direct way to determine if there is any tape in the reader or if the reader is on-line. In these situations an interrupt will occur immediately and an error indicated. As will be seen, the interrupt routine ignores errors if the state is *WAITING*, but if a good character comes in while in the *WAITING* state the interrupt routine sets the state to *READING*. Thus *open* loops until the state changes, meanwhile sleeping on the "lightning bolt" *lbolt*. If it did not sleep at all, it would prevent any other process from running until the reader came on-line; if it depended on the interrupt routine to wake it up, the effect would be the same, since the error interrupt is almost instantaneous.

The open-write case is much simpler: the punch is enabled and a 100-character leader is punched using *pleader*.

The *close* routine is also simple: if the reader was open, any uncollected characters are flushed, the interrupt is turned off, and the state is set to *CLOSED*. In the write case a 100-character trailer is punched. The routine has a bug in that if both the reader and punch are open *close* will be called only once, so that either the leftover characters are flushed or the trailer is punched, but not both. It is hard to see how to fix this problem except by making the reader and punch separate devices.

The *pread* routine tries to pick up characters from the input queue and passes them back until the user's read call is satisfied. If there are no characters it checks whether the state has gone to *EOF*, which means that the interrupt routine detected an error in the *READ* state (assumed to indicate the end of the tape). If so, *pread* returns; either during this call or the next one no characters will be passed back, indicating an end-of-file. If the state is still *READING* the routine enables another character by fiddling the device's reader control register, provided it is not active, and goes to sleep.

When a reader interrupt occurs and the state is *WAITING*, and the device's error bit is set, the interrupt is ignored; if there is no error the state is set to *READING*, as indicated in the discussion of *pread*. If the state is *READING* and there is an error, the state is set to *EOF*; it is assumed that the error represents the end of the tape. If there is no error, the character is picked up and stored in the input queue. Then, provided the number of characters already in the queue is less than the high-water mark *PCIHWAT*, the reader is enabled again to read another character. This strategy keeps the tape moving without flooding the input queue with unread characters. Finally, the top half is awakened.

Looking again at *pread*, notice that the priority level is raised by *spl4()* to prevent interrupts during the loop. This is done because of the possibility that the input queue is empty, and just after the EOF test is made an error interrupt occurs because the tape runs out. Then *sleep* will be called with no possibility of a *wakeup*. In general the processor priority should be raised when a routine is about to sleep awaiting some condition where the presence of the condition, and the consequent *wakeup*, is indicated by an interrupt. The danger is that the interrupt might occur between the test for the condition and the call to *sleep*, so that the *wakeup* apparently never happens.

At the same time it is a bad idea to raise the processor priority level for an extended period of time, since devices with real-time requirements may be shut out so long as to cause an error. The *pread* routine is perhaps overzealous in this respect, although since most devices have a priority level higher than 4 this difficulty is not very important.

The *pcwrite* routine simply gets characters from the user and passes them to *pcoutput*, which is separated out so that *pcleader* can call it also. *Pcoutput* checks for errors (like out-of-tape) and if none are present makes sure that the number of characters in the output queue does not exceed *PCOHWAT*; if it does, *sleep* is called. Then the character is placed on the output queue. There is a small bug here in that *pcoutput* does not check that the character was successfully put on the queue (all character-queue space might be empty); perhaps in this case it might be a good idea to sleep on the lightning-bolt until things quiet down. Finally *pcstart* is called, which checks to see if the punch is currently busy, and if not starts the punching of the first character on the output queue.

When punch interrupts occur, *pcpint* is called; it starts the punching of the next character on the output queue, and if the number of characters remaining on the queue is less than the low-water mark *PCOLWAT* it wakes up the write routine, which is presumably waiting for the queue to empty.

The Block-device Interface

Handling of block devices is mediated by a collection of routines which manage a set of buffers containing the images of blocks of data on the various devices. The most important purpose of these routines is to assure that several processes which access the same block of the same device in multiprogrammed fashion maintain a consistent view of the data in the block. A secondary but still important purpose is to increase the efficiency of the system by keeping in-core copies of blocks which are being accessed frequently. The main data base for this mechanism is the table of buffers *buf*. Each buffer header contains a pair of pointers (*b_forw*, *b_back*) which maintain a doubly-linked list of the buffers associated with a particular block device, and a pair of pointers (*av_forw*, *av_back*) which generally maintain a doubly-linked list of blocks which are "free," that is, eligible to be reallocated for another transaction. Buffers which have I/O in progress or are busy for other purposes do not appear in this list. The buffer header also contains the device and block number to which the buffer refers, and a pointer to the actual storage associated with the buffer. There is a word count which is the negative of the number of words to be transferred to or from the buffer; there is also an error byte and a residual word count used to communicate information from an I/O routine to its caller. Finally, there is a flag word with bits indicating the status of the buffer. These flags will be discussed below.

Six routines constitute the most important part of the interface with the rest of the system. Given a device and block number, both *bread* and *getblk* return a pointer to a buffer header for the block; the

difference is that *bread* is guaranteed to return a buffer actually containing the current data for the block, while *getblk* returns a buffer which contains the data in the block only if it is already in core (whether it is or not is indicated by the *B_DONE* bit; see below). In either case the buffer, and the corresponding device block, is made "busy," so that other processes referring to it are obliged to wait until it becomes free. *Getblk* is used, for example, when a block is about to be totally rewritten, so that its previous contents are not useful; still, no other process can be allowed to refer to the block until the new data is placed into it.

Given a pointer to a buffer, the *brelse* routine makes the buffer again available to other processes. It is called, for example, after data has been extracted following a *bread*. There are three subtly-different write routines, all of which take a buffer pointer as argument, and all of which logically release the buffer for use by others and place it on the free list. *Bwrite* puts the buffer on the appropriate device queue, waits for the write to be done, and sets the user's error flag if required.

Bawrite places the buffer on the device's queue, but does not wait for completion, so that errors cannot be reflected directly to the user. *Bdwrite* does not start any I/O operation at all, but merely marks the buffer so that if it happens to be grabbed from the free list to contain data from some other block, the data in it will first be written out.

Bwrite is used when one wants to be sure that I/O takes place correctly, and that errors are reflected to the proper user; it is used, for example, when updating i-nodes. *Bawrite* is useful when more efficiency is desired (because no wait is required for I/O to finish) but when it is reasonably certain that the write is really required. *Bdwrite* is used when there is doubt that the write is needed at the moment. For example, *bdwrite* is called when the last byte of a *write* system call falls short of the end of a block, on the assumption that another *write* will be given soon which will re-use the same block. On the other hand, as the end of a block is passed, *bawrite* is called, since probably the block will not be accessed again soon and one might as well start the writing process as soon as possible.

In any event, notice that the routines *getblk* and *bread* dedicate the given block exclusively to the use of the caller, and make others wait, while one of *brelse*, *bwrite*, *bawrite*, or *bdwrite* must eventually be called to free the block for use by others.

As mentioned, each buffer header contains a flag word which indicates the status of the buffer. Since they provide one important channel for information between the drivers and the block I/O system, it is important to understand them.

B_READ

This bit is set when the buffer is handed to the device strategy routine (see below) to indicate a read operation. The symbol *B_WRITE* is defined as 0 and does not define a flag; it is provided as a mnemonic convenience to callers of routines like *swap* which have a separate argument which indicates read or write.

B_DONE

This bit is set to 0 when a block is handed to the the device strategy routine and is turned on when the operation completes, whether normally as the result of an error. It is also used as part of the return argument of *getblk* to indicate if 1 that the returned buffer actually contains the data in the requested block.

B_ERROR

This bit may be set to 1 when *B_DONE* is set to indicate that an I/O or other error occurred. If it is set the *b_error* byte of the buffer header may contain an error code if it is non-zero. If *b_error* is 0 the nature of the error is not specified. Actually no driver at present sets *b_error*; the latter is provided for a future improvement whereby a more detailed error-reporting scheme may be implemented.

B_BUSY

This bit indicates that the buffer header is not on the free list, i.e. is dedicated to someone's exclusive use. The buffer still remains attached to the list of blocks associated with its device, however. When *getblk* (or *bread*, which calls it) searches the buffer list for a given device and finds the requested block

with this bit on, it sleeps until the bit clears.

B_WANTED

This flag is used in conjunction with the *B_BUSY* bit. Before sleeping as described just above, *getblk* sets this flag. Conversely, when the block is freed and the busy bit goes down (in *brelease*) a *wakeup* is given for the block header whenever *B_WANTED* is on. This strategem avoids the overhead of having to call *wakeup* every time a buffer is freed on the chance that someone might want it.

B_ASYNC

This bit is set by *bawrite* to indicate to the appropriate device driver that the buffer should be released when the write has been finished, usually at interrupt time. The difference between *bwrite* and *bawrite* is that the former starts I/O, waits until it is done, and frees the buffer. The latter merely sets this bit and starts I/O. The bit indicates that *relse* should be called for the buffer on completion.

B_DELWRI

This bit is set by *bdwrite* before releasing the buffer. When *getblk*, while searching for a free block, discovers the bit is 1 in a buffer it would otherwise grab, it causes the block to be written out before reusing it.

B_XMEM

This is actually a mask for the pair of bits which contain the high-order two bits of the physical address of the origin of the buffer; these bits are an extension of the 16 address bits elsewhere in the buffer header.

B_RELOC

This bit is currently unused; it previously indicated that a system-wide relocation constant was to be added to the buffer address. It was needed during a period when addresses of data in the system (including the buffers) were mapped by the relocation hardware to a physical address differing from its virtual address.

Block Device Drivers

The *bdevsw* table contains the names of the interface routines and that of a table for each block device.

Just as for character devices, block device drivers may supply an *open* and a *close* routine called respectively on each open and on the final close of the device. Instead of separate read and write routines, each block device driver has a *strategy* routine which is called with a pointer to a buffer header as argument. As discussed, the buffer header contains a read/write flag, the core address (including extended-memory bits), the block number, a (negative) word count, and the major and minor device number. The rôle of the strategy routine is to carry out the operation as requested by the information on the buffer header. When the transaction is complete the *B_DONE* (and possibly the *B_ERROR*) bits should be set. Then if the *B_ASYNC* bit is set, *brelease* should be called; otherwise, *wakeup*. In cases where the device is capable, under error-free operation, of transferring fewer words than requested, the device's word-count register should be placed in the residual count slot of the buffer header; otherwise, the residual count should be set to 0. This particular mechanism is really for the benefit of the magtape driver; when reading this device records shorter than requested are quite normal, and the user should be told the actual length of the record. [However the mechanism has not been integrated into normal I/O even on magtape and is used only in "raw" I/O as discussed below.]

Notice that although the most usual argument to the strategy routines is a genuine buffer header allocated as discussed above, all that is actually required is that the argument be a pointer to a place containing the appropriate information. For example the *swap* routine, which manages movement of core images to and from the swapping device, uses the strategy routine for this device. Care has to be taken that no extraneous bits get turned on in the flag word.

The device's table specified by *bdevsw* has a byte to contain an active flag and an error count, a pair of links which constitute the head of the chain of buffers for the device (*b_forw*, *b_back*), and a first and last

pointer for a device queue. Of these things, all are used solely by the device driver itself except for the buffer-chain pointers. Typically the flag encodes the state of the device, and is used at a minimum to indicate that the device is currently engaged in transferring information and no new command should be issued. The error count is useful for counting retries when errors occur. The device queue is used to remember stacked requests; in the simplest case it may be maintained as a first-in first-out list. Since buffers which have been handed over to the strategy routines are never on the list of free buffers, the pointers in the buffer which maintain the free list (*av_forw*, *av_back*) are also used to contain the pointers which maintain the device queues.

A couple of routines are provided which are useful to block device drivers.

iodone(bp)

given an pointer to a buffer header, arranges that it be released or awakened as appropriate to the situation when the strategy module has finished with the buffer whether normally or after an error. (In the latter case the *B_ERROR* bit has presumably been set.)

When the device conforms to some rather loose standards adhered to by certain DEC hardware, the routine

devstart(bp, devloc, devblk, hbcom)

is useful. Here *bp* is the address of the buffer header, *devloc* is the address of the slot in the device registers which accepts a perhaps-encoded device block number, *devblk* is the block number, and *hbcom* is a quantity to be stored in the high byte of the device's command register. It is understood, when using this routine, that the device registers are laid out in the order

command register
word count
core address
block (or track or sector)

where the address of the last corresponds to *devloc*. Moreover, the device should correspond to the RP, RK, and RF devices with respect to its layout of extended-memory bits and structure of read and write commands.

The routine

geterror(bp)

can be used to examine the error bit in a buffer header and arrange that any error indication found therein is reflected to the user. It may be called only in the non-interrupt part of a driver when I/O has completed (*B_DONE* has been set).

An example

The RF disk driver is worth studying as the simplest example of a block I/O device. Its *strategy* routine checks to see if the requested block lies beyond the end of the device; the size of the disk, in this instance, is indicated by the minor device number. If the request is plausible, the buffer is placed at the end of the device queue, and if the disk is not running, *rfstart* is called.

Rfstart merely returns if there is nothing to do, but otherwise sets the device-active flag, loads the address extension register, and calls *devstart* to perform the remaining tasks attendant on beginning a data transfer.

When a completion or error interrupt occurs, *rfintr* is called. If an error is indicated, and if the error count has not exceeded 10, the same transaction is reattempted; otherwise the error bit is set. If there was no error or if 10 failing transfers have been issued the queue is advanced and *rfstart* is called to begin another transaction.

Raw Block-device I/O

A scheme has been set up whereby block device drivers may provide the ability to transfer information directly between the user's core image and the device without the use of buffers and in blocks as large as the caller requests. The method involves setting up a character-type special file corresponding to the raw

device and providing *read* and *write* routines which set up what is usually a private, non-shared buffer header with the appropriate information and call the device's strategy routine. If desired, separate *open* and *close* routines may be provided but this is usually unnecessary. A special-function routine might come in handy, especially for magtape.


A great deal of work has to be done to generate the "appropriate information" to put in the argument buffer for the strategy module; the worst part is to map relocated user addresses to physical addresses. Most of this work is done by


physio(strat, bp, dev, rw)

whose arguments are the name of the strategy routine *strat*, the buffer pointer *bp*, the device number *dev*, and a read-write flag *rw* whose values is either *B_READ* or *B_WRITE*. *Physio* makes sure that the user's base address and count are even (because most devices work in words) and that the core area affected is contiguous in physical space; it delays until the buffer is not busy, and makes it busy while the operation is in progress, and it sets up user error return information.

The magtape driver is the only one which as of this writing provides a raw I/O capability; given *physio*, the work involved is trivial, and amounts to passing back to the user information on the length of the record read or written. (There is some funniness because the magtape, uniquely among DEC devices, works in bytes, not words.) Putting in raw I/O for disks should be equally trivial except that the disk address has to be carefully checked to make sure it does not overflow from one logical device to another on which the caller may not have write permission.







On the Security of UNIX

Dennis M. Ritchie

Bell Laboratories
Murray Hill, New Jersey 07974

Recently there has been much interest in the security aspects of operating systems and software. At issue is the ability to prevent undesired disclosure of information, destruction of information, and harm to the functioning of the system. This paper discusses the degree of security which can be provided under the UNIX system and offers a number of hints on how to improve security.

The first fact to face is that UNIX was not developed with security, in any realistic sense, in mind; this fact alone guarantees a vast number of holes. (Actually the same statement can be made with respect to most systems.) The area of security in which UNIX is theoretically weakest is in protecting against crashing or at least crippling the operation of the system. The problem here is not mainly in uncritical acceptance of bad parameters to system calls— there may be bugs in this area, but none are known— but rather in the lack of any checks for excessive consumption of resources. Most notably, there is no limit on the amount of disk storage used, either in total space allocated or in the number of files or directories. Here is a particularly ghastly shell sequence guaranteed to stop the system:

```
: loop
mkdir x
chdir x
goto loop
```

Either a panic will occur because all the i-nodes on the device are used up or all the disk blocks will be consumed, thus preventing anyone from writing files on the device.

Processes are another resource on which the only limit is total exhaustion. For example, the sequence

```
command&
command&
command&
...
```

if continued long enough will use up all the slots in the system's process table and prevent anyone from executing any commands. Alternatively, if the commands use much core, swap space may run out, causing a panic. Incidentally, because of the implementation of process termination, the above sequence is effective in stopping the system no matter how short a time it takes each command to terminate. (The process-table slot is not freed until the terminated process is waited for; if no commands without "&" are executed, the Shell never executes a "wait.")

It should be evident that unbounded consumption of disk space, files, swap space, and processes can easily occur accidentally in malfunctioning programs as well as at command level. In fact UNIX is essentially defenseless against this kind of abuse, nor is there any easy fix. The best that can be said is that it is generally fairly easy to detect what has happened when disaster strikes, to identify the user responsible, and take appropriate action. In practice, we have found that difficulties in this area are rather rare, but we have not been faced with malicious users, and enjoy a fairly generous supply of resources which have served to cushion us against

accidental overconsumption.

The picture is considerably brighter in the area of protection of information from unauthorized perusal and destruction. Here the degree of security seems (almost) adequate theoretically, and the problems lie more in the necessity for care in the actual use of the system.

Each UNIX file has associated with it eleven bits of protection information together with a user identification number and a user-group identification number (UID and GID). Nine of the protection bits are used to specify independently permission to read, to write, and to execute the file to the user himself, to members of the user's group, and to all other users. Each process generated by or for a user has associated with it an effective UID and a real UID, and an effective and real GID. When an attempt is made to access the file for reading, writing, or execution, the user process's effective UID is compared against the file's UID; if a match is obtained, access is granted provided the read, write, or execute bit respectively for the user himself is present. If the UID for the file and for the process fail to match, but the GID's do match, the group bits are used; if the GID's do not match, the bits for other users are tested. The last two bits of each file's protection information, called the set-UID and set-GID bits, are used only when the file is executed as a program. If, in this case, the set-UID bit is on for the file, the effective UID for the process is changed to the UID associated with the file; the change persists until the process terminates or until the UID changed again by another execution of a set-UID file. Similarly the effective group ID of a process is changed to the GID associated with a file when that file is executed and has the set-GID bit set. The real UID and GID of a process do not change when any file is executed, but only as the result of a privileged system call.

The basic notion of the set-UID and set-GID bits is that one may write a program which is executable by others and which maintains files accessible to others only by that program. The classical example is the game-playing program which maintains records of the scores of its players. The program itself has to read and write the score file, but no one but the game's sponsor can be allowed unrestricted access to the file lest they manipulate the game to their own advantage. The solution is to turn on the set-UID bit of the game program. When, and only when, it is invoked by players of the game, it may update the score file ordinary programs executed by others cannot access the score.

There are a number of special cases involved in determining access permissions. Since executing a directory as a program is a meaningless operation, the execute-permission bit, for directories, is taken instead to mean permission to search the directory for a given file during the scanning of a path name; thus if a directory has execute permission but no read permission for a given user, he may access files with known names in the directory, but may not read (list) the entire contents of the directory. Write permission on a directory is interpreted to mean that the user may create and delete files in that directory; it is impossible for any user to write directly into any directory.

Another, and from the point of view of security, much more serious special case is that there is a "super user" who is able to read any file and write any non-directory. The super-user is also able to change the protection mode and the owner UID and GID of any file and to invoke privileged system calls. It must be recognized that the mere notion of a super-user is a theoretical, and usually practical, blemish on any protection scheme.

The first necessity for a secure system is of course arranging that all files and directories have the proper protection modes. Unfortunately, UNIX software is exceedingly permissive in this regard; essentially all commands create files readable and writable by everyone. This means that more or less continuous attention must be paid to adjusting modes properly. If one wants to keep one's files completely secret, it is possible to remove all permissions from the directory in which they live, which is easy and effective; but if it is desired to give general read permission while preventing writing, things are more complicated. The main problem is that write permission in a directory means precisely that; it has nothing to do with write permission

for a file in that directory. Thus a writeable file in a read-only directory may be changed, or even truncated, though not removed. This fact is perfectly logical, though in this case unfortunate. A case can be made for requiring write permission for the directory of a file as well as for the file itself before allowing writing. (This possibility is more complicated than it seems at first; the system has to allow users to change their own directories while forbidding them to change the user-directory directory.)

A situation converse to the above-discussed difficulty is also present— it is possible to delete a file if one has write permission for its directory independently of any permissions for the file. This problem is related more to self-protection than protection from others. It is largely mitigated by the fact that the two major commands which delete named files (*mv* and *rm*) ask confirmation before deleting unwritable files.

It follows from this discussion that to maintain both data privacy and data integrity, it is necessary, and largely sufficient, to make one's directory inaccessible to others. The lack of sufficiency could follow from the existence of set-UID programs created by the user and the possibility of total breach of system security in one of the ways discussed below (or one of the ways not discussed below).

Needless to say, the system administrators must be at least as careful as their most demanding user to place the correct protection mode on the files under their control. In particular, it is necessary that special files be protected from writing, and probably reading, by ordinary users when they store sensitive files belonging to other users. It is easy to write programs that examine and change files by accessing the device on which the files live.


On the issue of password security, UNIX is probably better than most systems. Passwords are stored in an encrypted form which, in the absence of serious attention from specialists in the field, appears reasonably secure, provided its limitations are understood. Since both the encryption algorithm and the encrypted passwords are available, exhaustive enumeration of potential passwords is feasible up to a point. As a practical test of the possibilities in this area, 67 encrypted passwords were collected from 10 UNIX installations. These were tested against all five-letter combinations, all combinations of letters and digits of length four or less, and all words in Webster's Second unabridged dictionary; 60 of the 67 passwords were found. The whole process took about 12 hours of machine time. This experience suggests that passwords should be at least six characters long and randomly chosen from an alphabet which includes digits and special characters.


Of course there also exist feasible non-cryptanalytic ways of finding out passwords. For example: write a program which types out "login:" on the typewriter and copies whatever is typed to a file of your own. Then invoke the command and go away until the victim arrives. (It is this kind of possibility that makes it evident that UNIX was not designed to be secure.)

The set-UID (set-GID) notion must be used carefully if any security is to be maintained. The first thing to keep in mind is that a writable set-UID file can have another program copied onto it. For example, if the super-user (*su*) command is writable, anyone can copy the shell onto it and get a password-free version of *su*. A more subtle problem can come from set-UID programs which are not sufficiently careful of what is fed into them. In some systems, for example, the *mail* command is set-UID and owned by the super-user. The notion is that one should be able to send mail to anyone even if they want to protect their directories from writing. The trouble is that *mail* is rather dumb: anyone can mail someone else's private file to himself. Much more serious, is the following scenario: make a file with a line like one in the password file which allows one to log in as the super-user. Then make a link named ".mail" to the password file in some writeable directory on the same device as the password file (say /tmp). Finally mail the bogus login line to /tmp/.mail. You can then login as the super-user, clean up the incriminating evidence, and have your will.

The fact that users can mount their own disks and tapes as file systems can be another

way of gaining super-user status. Once a disk pack is mounted, the system believes what is on it. Thus one can take a blank disk pack, put on it anything desired, and mount it. There are obvious and unfortunate consequences. For example: a mounted disk with garbage on it will crash the system; one of the files on the mounted disk can easily be a password-free version of *su*; other files can be unprotected entries for special files. The only easy fix for this problem is to forbid the use of *mount* to unprivileged users. A partial solution, not so restrictive, would be to have the *mount* command examine the special file for bad data, set-UID programs owned by others, and accessible special files, and balk at unprivileged invokers.





UNIX Assembler Reference Manual

Dennis M. Ritchie

*Bell Laboratories
Murray Hill, New Jersey*

0. Introduction

This document describes the usage and input syntax of the UNIX PDP-11 assembler *as*. The details of the PDP-11 are not described; consult the DEC documents "PDP-11/20 Handbook" and "PDP-11/45 Handbook."

The input syntax of the UNIX assembler is generally similar to that of the DEC assembler PAL-11R, although its internal workings and output format are unrelated. It may be useful to read the publication DEC-11-ASDB-D, which describes PAL-11R, although naturally one must use care in assuming that its rules apply to *as*.

As is a rather ordinary two-pass assembler without macro capabilities. It produces an output file which contains relocation information and a complete symbol table; thus the output is acceptable to the UNIX link-editor *ld*, which may be used to combine the outputs of several assembler runs and to obtain object programs from libraries. The output format has been designed so that if a program contains no unresolved references to external symbols, it is executable without further processing.

1. Usage

as is used as follows:

```
as [ - ] file, ...
```

If the optional "-" argument is given, all undefined symbols in the current assembly will be made undefined-external. See the `.globl` directive below.

The other arguments name files which are concatenated and assembled. Thus programs may be written in several pieces and assembled together.

The output of the assembler is placed on the file *a.out* in the current directory. If there were no unresolved external references, and no errors detected, *a.out* is made executable; otherwise, if it is produced at all, it is made non-executable.

2. Lexical conventions

Assembler tokens include identifiers (alternatively, "symbols" or "names"), temporary symbols, constants, and operators.

2.1 Identifiers

An identifier consists of a sequence of alphanumeric characters (including period ".", underscore "_", and tilde "~" as alphanumeric) of which the first may not be numeric. Only the first eight characters are significant. When a name begins with a tilde, the tilde is discarded and that occurrence of the identifier generates a unique entry in the symbol table which can match no other occurrence of the identifier. This feature is used by the C compiler to place names of local variables in the output sym-

bol table without having to worry about making them unique.

2.2 Temporary symbols

A temporary symbol consists of a digit followed by “f” or “b”. Temporary symbols are discussed fully in §5.1.

2.3 Constants

An octal constant consists of a sequence of digits; “8” and “9” are taken to have octal value 10 and 11. The constant is truncated to 16 bits and interpreted in two’s complement notation.

A decimal constant consists of a sequence of digits terminated by a decimal point “.”. The magnitude of the constant should be representable in 15 bits; i.e., be less than 32,768.

A single-character constant consists of a single quote “'” followed by an ASCII character not a new-line. Certain dual-character escape sequences are acceptable in place of the ASCII character to represent new-line and other non-graphics (see *String statements*, §5.5). The constant’s value has the code for the given character in the least significant byte of the word and is null-padded on the left.

A double-character constant consists of a double quote “”” followed by a pair of ASCII characters not including new-line. Certain dual-character escape sequences are acceptable in place of either of the ASCII characters to represent new-line and other non-graphics (see *String statements*, §5.5). The constant’s value has the code for the first given character in the least significant byte and that for the second character in the most significant byte.

2.4 Operators

There are several single- and double-character operators; see §6.

2.5 Blanks

Blank and tab characters may be interspersed freely between tokens, but may not be used within tokens (except character constants). A blank or tab is required to separate adjacent identifiers or constants not otherwise separated.

2.6 Comments

The character “/” introduces a comment, which extends through the end of the line on which it appears. Comments are ignored by the assembler.

3. Segments

Assembled code and data fall into three segments: the text segment, the data segment, and the bss segment. The text segment is the one in which the assembler begins, and it is the one into which instructions are typically placed. The UNIX system will, if desired, enforce the purity of the text segment of programs by trapping write operations into it. Object programs produced by the assembler must be processed by the link-editor *ld* (using its “-n” flag) if the text segment is to be write-protected. A single copy of the text segment is shared among all processes executing such a program.

The data segment is available for placing data or instructions which will be modified during execution. Anything which may go in the text segment may be put into the data segment. In programs with write-protected, sharable text segments, data segment contains the initialized but variable parts of a program. If the text segment is not pure, the data segment begins immediately after the text segment; if the text segment is pure, the data segment begins at the lowest 8K byte boundary after the text segment.

The bss segment may not contain any explicitly initialized code or data. The length of the bss segment (like that of text or data) is determined by the high-water mark of the location counter within it. The bss segment is actually an extension of the data segment and begins immediately after it. At the start of execution of a program, the bss segment is set to 0. Typically the bss segment is set up by statements exemplified by

```
lab: . = .+10
```


The advantage in using the bss segment for storage that starts off empty is that the initialization information need not be stored in the output file. See also *Location counter* and *Assignment statements* below.

4. The location counter

One special symbol, ".", is the location counter. Its value at any time is the offset within the appropriate segment of the start of the statement in which it appears. The location counter may be assigned to, with the restriction that the current segment may not change; furthermore, the value of "." may not decrease. If the effect of the assignment is to increase the value of ".", the required number of null bytes are generated (but see *Segments* above).

5. Statements

A source program is composed of a sequence of *statements*. Statements are separated either by newlines or by semicolons. There are five kinds of statements: null statements, expression statements, assignment statements, string statements, and keyword statements.

Any kind of statement may be preceded by one or more labels.

5.1 Labels

There are two kinds of label: name labels and numeric labels. A name label consists of a name followed by a colon (:). The effect of a name label is to assign the current value and type of the location counter "." to the name. An error is indicated in pass 1 if the name is already defined; an error is indicated in pass 2 if the "." value assigned changes the definition of the label.

A numeric label consists of a digit 0 to 9 followed by a colon (:). Such a label serves to define temporary symbols of the form "nb" and "nf", where *n* is the digit of the label. As in the case of name labels, a numeric label assigns the current value and type of "." to the temporary symbol. However, several numeric labels with the same digit may be used within the same assembly. References of the form "nf" refer to the first numeric label "n:" forward from the reference; "nb" symbols refer to the first "n:" label backward from the reference. This sort of temporary label was introduced by Knuth [*The Art of Computer Programming, Vol I: Fundamental Algorithms*]. Such labels tend to conserve both the symbol table space of the assembler and the inventive powers of the programmer.

5.2 Null statements

A null statement is an empty statement (which may, however, have labels). A null statement is ignored by the assembler. Common examples of null statements are empty lines or lines containing only a label.

5.3 Expression statements

An expression statement consists of an arithmetic expression not beginning with a keyword. The assembler computes its (16-bit) value and places it in the output stream, together with the appropriate relocation bits.

5.4 Assignment statements

An assignment statement consists of an identifier, an equals sign (=), and an expression. The value and type of the expression are assigned to the identifier. It is not required that the type or value be the same in pass 2 as in pass 1, nor is it an error to redefine any symbol by assignment.

Any external attribute of the expression is lost across an assignment. This means that it is not possible to declare a global symbol by assigning to it, and that it is impossible to define a symbol to be offset from a non-locally defined global symbol.

As mentioned, it is permissible to assign to the location counter ".". It is required, however, that the type of the expression assigned be of the same type as ".", and it is forbidden to decrease the value of ".". In practice, the most common assignment to "." has the form ". = . + *n*" for some number *n*;

this has the effect of generating *n* null bytes.

5.5 String statements

A string statement generates a sequence of bytes containing ASCII characters. A string statement consists of a left string quote "<" followed by a sequence of ASCII characters not including newline, followed by a right string quote ">". Any of the ASCII characters may be replaced by a two-character escape sequence to represent certain non-graphic characters, as follows:

\n	NL	(012)
\t	HT	(011)
\e	EOT	(004)
\0	NUL	(000)
\r	CR	(015)
\a	ACK	(006)
\p	PFX	(033)
\\	\	
\>	>	

The last two are included so that the escape character and the right string quote may be represented. The same escape sequences may also be used within single- and double-character constants (see §2.3 above).

5.6 Keyword statements

Keyword statements are numerically the most common type, since most machine instructions are of this sort. A keyword statement begins with one of the many predefined keywords of the assembler; the syntax of the remainder depends on the keyword. All the keywords are listed below with the syntax they require.

6. Expressions

An expression is a sequence of symbols representing a value. Its constituents are identifiers, constants, temporary symbols, operators, and brackets. Each expression has a type.

All operators in expressions are fundamentally binary in nature; if an operand is missing on the left, a 0 of absolute type is assumed. Arithmetic is two's complement and has 16 bits of precision. All operators have equal precedence, and expressions are evaluated strictly left to right except for the effect of brackets.

6.1 Expression operators

The operators are:

(blank)	when there is no operator between operands, the effect is exactly the same as if a "+" had appeared.
+	addition
-	subtraction
*	multiplication
∕	division (note that plain "/" starts a comment)
&	bitwise and
	bitwise or
>>	logical right shift
<<	logical left shift
..	modulo
'	<i>a!b</i> is <i>a or (not b)</i> ; i.e., the or of the first operand and the one's complement of the

second; most common use is as a unary.

result has the value of first operand and the type of the second; most often used to define new machine instructions with syntax identical to existing instructions.

Expressions may be grouped by use of square brackets "[]". (Round parentheses are reserved for address modes.)

6.2 Types

The assembler deals with a number of types of expressions. Most types are attached to keywords and used to select the routine which treats that keyword. The types likely to be met explicitly are:

undefined

Upon first encounter, each symbol is undefined. It may become undefined if it is assigned an undefined expression. It is an error to attempt to assemble an undefined expression in pass 2; in pass 1, it is not (except that certain keywords require operands which are not undefined).

undefined external

A symbol which is declared `.globl` but not defined in the current assembly is an undefined external. If such a symbol is declared, the link editor `ld` must be used to load the assembler's output with another routine that defines the undefined reference.

absolute

An absolute symbol is one defined ultimately from a constant. Its value is unaffected by any possible future applications of the link-editor to the output file.

text

The value of a text symbol is measured with respect to the beginning of the text segment of the program. If the assembler output is link-edited, its text symbols may change in value since the program need not be the first in the link editor's output. Most text symbols are defined by appearing as labels. At the start of an assembly, the value of "." is text 0.

data

The value of a data symbol is measured with respect to the origin of the data segment of a program. Like text symbols, the value of a data symbol may change during a subsequent link-editor run since previously loaded programs may have data segments. After the first `.data` statement, the value of "." is data 0.

bss

The value of a bss symbol is measured from the beginning of the bss segment of a program. Like text and data symbols, the value of a bss symbol may change during a subsequent link-editor run, since previously loaded programs may have bss segments. After the first `.bss` statement, the value of "." is bss 0.

external absolute, text, data, or bss

symbols declared `.globl` but defined within an assembly as absolute, text, data, or bss symbols may be used exactly as if they were not declared `.globl`; however, their value and type are available to the link editor so that the program may be loaded with others that reference these symbols.

register

The symbols

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

are predefined as register symbols. Either they or symbols defined from them must be used to refer to the six general-purpose, six floating-point, and the 2 special-purpose machine registers. The behavior of the floating register names is identical to that of the

corresponding general register names; the former are provided as a mnemonic aid.

other types

Each keyword known to the assembler has a type which is used to select the routine which processes the associated keyword statement. The behavior of such symbols when not used as keywords is the same as if they were absolute.

6.3 Type propagation in expressions

When operands are combined by expression operators, the result has a type which depends on the types of the operands and on the operator. The rules involved are complex to state but were intended to be sensible and predictable. For purposes of expression evaluation the important types are

- undefined
- absolute
- text
- data
- bss
- undefined external
- other

The combination rules are then: If one of the operands is undefined, the result is undefined. If both operands are absolute, the result is absolute. If an absolute is combined with one of the "other types" mentioned above, or with a register expression, the result has the register or other type. As a consequence, one can refer to r3 as "r0+3". If two operands of "other type" are combined, the result has the numerically larger type (not that this fact is very useful, since the values are not made public). An "other type" combined with an explicitly discussed type other than absolute acts like an absolute.

Further rules applying to particular operators are:

- + If one operand is text-, data-, or bss-segment relocatable, or is an undefined external, the result has the postulated type and the other operand must be absolute.
 - If the first operand is a relocatable text-, data-, or bss-segment symbol, the second operand may be absolute (in which case the result has the type of the first operand); or the second operand may have the same type as the first (in which case the result is absolute). If the first operand is external undefined, the second must be absolute. All other combinations are illegal.
- This operator follows no other rule than that the result has the value of the first operand and the type of the second.

others It is illegal to apply these operators to any but absolute symbols.

7. Pseudo-operations

The keywords listed below introduce statements which generate data in unusual forms or influence the later operations of the assembler. The metanotation

[stuff] ...

means that 0 or more instances of the given stuff may appear. Also, boldface tokens are literals, italic words are substitutable.

7.1 **.byte** *expression* [, *expression*] ...

The *expressions* in the comma-separated list are truncated to 8 bits and assembled in successive bytes. The expressions must be absolute. This statement and the string statement above are the only ones which assemble data one byte at a time.

7.2 **even**

If the location counter "." is odd, it is advanced by one so the next statement will be assembled at a

word boundary.

7.3 *.if expression*

The *expression* must be absolute and defined in pass 1. If its value is nonzero, the *.if* is ignored; if zero, the statements between the *.if* and the matching *.endif* (below) are ignored. *.if* may be nested. The effect of *.if* cannot extend beyond the end of the input file in which it appears. (The statements are not totally ignored, in the following sense: *.ifs* and *.endifs* are scanned for, and moreover all names are entered in the symbol table. Thus names occurring only inside an *.if* will show up as undefined if the symbol table is listed.)

7.4 *.endif*

This statement marks the end of a conditionally-assembled section of code. See *.if* above.

7.5 *.globl name [, name] ...*

This statement makes the *names* external. If they are otherwise defined (by assignment or appearance as a label) they act within the assembly exactly as if the *.globl* statement were not given; however, the link editor *ld* may be used to combine this routine with other routines that refer these symbols.

Conversely, if the given symbols are not defined within the current assembly, the link editor can combine the output of this assembly with that of others which define the symbols.

As discussed in §1, it is possible to force the assembler to make all otherwise undefined symbols external.

7.6 *.text*

7.7 *.data*

7.8 *.bss*

These three pseudo-operations cause the assembler to begin assembling into the text, data, or bss segment respectively. Assembly starts in the text segment. It is forbidden to assemble any code or data into the bss segment, but symbols may be defined and “.” moved about by assignment.

7.9 *.comm name , expression*

Provided the *name* is not defined elsewhere, this statement is equivalent to

```
globl name
name = expression ^ name
```

That is, the type of *name* is “undefined external”, and its value is *expression*. In fact the *name* behaves in the current assembly just like an undefined external. However, the link-editor *ld* has been special-cased so that all external symbols which are not otherwise defined, and which have a non-zero value, are defined to lie in the bss segment, and enough space is left after the symbol to hold *expression* bytes. All symbols which become defined in this way are located before all the explicitly defined bss-segment locations.

8. Machine instructions

Because of the rather complicated instruction and addressing structure of the PDP-11, the syntax of machine instruction statements is varied. Although the following sections give the syntax in detail, the 11/20 and 11/45 handbooks should be consulted on the semantics.

8.1 Sources and Destinations

The syntax of general source and destination addresses is the same. Each must have one of the following forms, where *reg* is a register symbol, and *expr* is any sort of expression:

syntax	words	mode
<i>reg</i>	0	0+ <i>reg</i>
(<i>reg</i>)+	0	2+ <i>reg</i>
-(<i>reg</i>)	0	4+ <i>reg</i>
<i>expr</i> (<i>reg</i>)	1	6+ <i>reg</i>
(<i>reg</i>)	0	1+ <i>reg</i>
* <i>reg</i>	0	1+ <i>reg</i>
*(<i>reg</i>)+	0	3+ <i>reg</i>
*-(<i>reg</i>)	0	5+ <i>reg</i>
*(<i>reg</i>)	1	7+ <i>reg</i>
* <i>expr</i> (<i>reg</i>)	1	7+ <i>reg</i>
<i>expr</i>	1	67
\$ <i>expr</i>	1	27
* <i>expr</i>	1	77
*\$ <i>expr</i>	1	37

The *words* column gives the number of address words generated; the *mode* column gives the octal address-mode number. The syntax of the address forms is identical to that in DEC assemblers, except that "*" has been substituted for "@" and "\$" for "#"; the UNIX typing conventions make "@" and "#" rather inconvenient.

Notice that mode "*reg" is identical to "(reg)"; that "*(reg)" generates an index word (namely, 0); and that addresses consisting of an unadorned expression are assembled as pc-relative references independent of the type of the expression. To force a non-relative reference, the form "\$expr" can be used, but notice that further indirection is impossible.

8.3 Simple machine instructions

The following instructions are defined as absolute symbols:

```
clc
clv
clz
cln
sec
sev
sez
sen
```

They therefore require no special syntax. The PDP-11 hardware allows more than one of the "clear" class, or alternatively more than one of the "set" class to be or-ed together; this may be expressed as follows:

```
clc | clv
```

8.4 Branch

The following instructions take an expression as operand. The expression must lie in the same segment as the reference, cannot be undefined-external, and its value cannot differ from the current location of "." by more than 254 bytes:

br	blos
bne	bvc
beq	bvs
bge	bhis
blt	bec (= bcc)
bgt	bcc
ble	blo
bpl	bcs
bmi	bes (= bcs)
bhi	

bes ("branch on error set") and **bec** ("branch on error clear") are intended to test the error bit returned by system calls (which is the c-bit).

8.5 Extended branch instructions

The following symbols are followed by an expression representing an address in the same segment as "...". If the target address is close enough, a branch-type instruction is generated; if the address is too far away, a **jmp** will be used.

jbr	jlos
jne	jvc
jeq	jvs
jge	jhis
jlt	jec
jgt	jcc
jle	jlo
jpl	jcs
jmi	jes
jhi	

jbr turns into a plain **jmp** if its target is too remote; the others (whose names are constructed by replacing the "b" in the branch instruction's name by "j") turn into the converse branch over a **jmp** to the target address.

8.6 Single operand instructions

The following symbols are names of single-operand machine instructions. The form of address expected is discussed in §8.1 above.

clr	sbeb
clrb	rør
com	rørb
comb	roi
inc	rolb
incb	asr
dec	asrb
decb	asl
neg	aslb
negb	jmp
adc	swab
adcb	tst
sbc	tstb

8.7 Double operand instructions

The following instructions take a general source and destination (§8.1), separated by a comma, as operands

mov
movb
cmp
cmpb
bit
bitb
bic
bicb
bis
bisb
add
sub

8.8 Miscellaneous instructions

The following instructions have more specialized syntax. Here *reg* is a register name, *src* and *dst* a general source or destination (§8.1), and *expr* is an expression:

jsr	<i>reg, dst</i>	
rts	<i>reg</i>	
sys	<i>expr</i>	
ash	<i>src, reg</i>	(or, als)
ashc	<i>src, reg</i>	(or, alsc)
mul	<i>src, reg</i>	(or, mpy)
div	<i>src, reg</i>	(or, dvd)
xor	<i>reg, dst</i>	
sxt	<i>dst</i>	
mark	<i>expr</i>	
sob	<i>reg, expr</i>	

sys is another name for the **trap** instruction. It is used to code system calls. Its operand is required to be expressible in 6 bits. The alternative forms for **ash**, **ashc**, **mul**, and **div** are provided to avoid conflict with EAE register names should they be needed.

The expression in **mark** must be expressible in six bits, and the expression in **sob** must be in the same segment as "...", must not be external-undefined, must be less than "...", and must be within 510 bytes of "...".

8.9 Floating-point unit instructions

The following floating-point operations are defined, with syntax as indicated:

cfcc		
setf		
setd		
seti		
setl		
clrf	<i>fdst</i>	
negf	<i>fdst</i>	
absf	<i>fdst</i>	
tstf	<i>fsrc</i>	
movf	<i>fsrc, freg</i>	(= ldf)
movf	<i>freg, fdst</i>	(= stf)
movif	<i>src, freg</i>	(= ldcif)
movfi	<i>freg, dst</i>	(= stcfi)
movof	<i>fsrc, freg</i>	(= ldcdf)
movfo	<i>freg, fdst</i>	(= stcfd)
movie	<i>src, freg</i>	(= ldexp)
movei	<i>freg, dst</i>	(= stexp)
addf	<i>fsrc, freg</i>	
subf	<i>fsrc, freg</i>	

mulf	<i>fsrc, freg</i>
divf	<i>fsrc, freg</i>
cmpf	<i>fsrc, freg</i>
modf	<i>fsrc, freg</i>
ldfps	<i>src</i>
stfps	<i>dst</i>
stst	<i>dst</i>

fsrc, *fdst*, and *freg* mean floating-point source, destination, and register respectively. Their syntax is identical to that for their non-floating counterparts, but note that only floating registers 0—3 can be a *freg*.

The names of several of the operations have been changed to bring out an analogy with certain fixed-point instructions. The only strange case is **movf**, which turns into either **stf** or **ldf** depending respectively on whether its first operand is or is not a register. Warning: **ldf** sets the floating condition codes, **stf** does not.

9. Other symbols

9.1 ..

The symbol “..” is the *relocation counter*. Just before each assembled word is placed in the output stream, the current value of this symbol is added to the word if the word refers to a text, data or bss segment location. If the output word is a pc-relative address word which refers to an absolute location, the value of “..” is subtracted.

Thus the value of “..” can be taken to mean the starting core location of the program. In UNIX systems with relocation hardware, the initial value of “..” is 0.

The value of “..” may be changed by assignment. Such a course of action is sometimes necessary, but the consequences should be carefully thought out. It is particularly ticklish to change “..” midway in an assembly or to do so in a program which will be treated by the loader, which has its own notions of “..”.

9.2 System calls

The following absolute symbols may be used to code calls to the UNIX system (see the **sys** instruction above).

break	nice
chdir	open
chmod	read
chown	seek
close	setuid
creat	signal
exec	stat
exit	stime
fork	stty
fstat	tell
getuid	time
gtty	umount
link	unlink
mkdir	wait
mdate	write
mount	

Warning: the **wait** system call is not the same as the **wait** instruction, which is not defined in the assembler.

10. 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 address 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 phase error— “.” different in pass 1 and 2
- R relocation error
- U undefined symbol
- X syntax error

PWB/UNIX Manual Page Macros

E. M. Piskorik

Bell Laboratories
Piscataway, New Jersey 08854

This memorandum is a reference for all who wish to prepare manual pages like those in the PWB/UNIX* User's Manual [1] or in the UNIX Programmer's Manual [2]. The descriptions of the available macros¹ are particularized for use in the PWB/UNIX environment, but they are useful as well for all who wish to prepare manual pages for other similar manuals. Described below are the available macros for both *nroff* and *troff* [3] output with examples and explanations of their use.

TITLE HEADING

Manual pages consist of a title heading, several subject headings, and indented text paragraphs. The title heading is defined by the `.th` macro, which takes three arguments² and is written as follows:

```
.th name section date
```

For example:

```
.th GSI I 11/10/77
```

The `.th` macro results in a heading consisting of a centered *date* with the *name* and *section* (in parentheses) appearing together in the left and right margins at the top of the page. This macro *must* appear before any other text, macro, or request so that the proper page heading is generated.

SUBJECT HEADINGS

The subject headings NAME, SYNOPSIS, and DESCRIPTION must appear on all manual pages. It is strongly recommended that the headings FILES, SEE ALSO, and, if appropriate, BUGS be included. Other headings such as the following are optional and may be included as the author sees fit:

```
DIAGNOSTICS
EXAMPLES
WARNINGS
RANGES AND LIMITATIONS
CAVEAT
```

The `.sh` macro is used for subject headings and takes the heading as its argument. The heading is printed at the left-hand margin. Text that follows is single spaced and indented five ens³ beneath the heading. Thus:

```
.sh NAME
text lines
:
:
```

produces:

```
NAME
text lines
:
:
```

* UNIX is a Trademark/Service Mark of the Bell System.

1 The original version of these macros was written by D. M. Ritchie and K. Thompson.

2 Any macro argument that contains blanks must be enclosed in double quotes, e. g., "string of words".

3 An *en* is the number of points (1 point = 1/72 of an inch) equal to half the current point size. In *nroff* an *en* is equal to a character width.

SIMPLE PARAGRAPHS

New paragraphs under any subject heading are produced by calling one of the `.s1`, `.s2`, and `.s3` macros. The `.s1` macro outputs a vertical space and makes sure that there is room for at least three more lines on the page. In *nroff*, both the `.s2` and `.s3` macros output a vertical space, while in *troff*, they output one half of a vertical space. The `.s3` macro differs from `.s2` in that `.s3` ensures that there is enough room on the page for three more lines, while `.s2` does not.

LABELED PARAGRAPHS

The `.lp` macro requires two arguments and is designed to produce a "tagged" list, where a tag or item appears on the first line in the first column and the explanatory text appears in the second column opposite the tag on as many lines as are necessary. The first argument to this macro is the amount of text indentation from the left margin to the text, and the second is the amount to the *left* of the text where the tag is to start. Thus, the command "`.lp 14 7`" places the text 14 ens from the left margin (i.e., 9 ens to the right of the *previous text*) and places the tag associated with the text 7 ens to the *left* of the text. The first input line after the `.lp` macro should contain the tag or item followed by a tab and the text. On output, the formatter automatically fills the lines of text. For example, the sequence:

```

.s3
.lp 14 7
+12→permits use of 12-pitch,
6 lines/inch text.
GSI terminals normally
allow only two combinations:
10-pitch, 6 lines/inch, or
12-pitch, 8 lines/inch.
.s3
.lp 14 7
\ -n→controls the size of the
half-line spacing.
A half-line is, by default, equal
to 4 vertical plot increments.
A 10-pitch line-feed requires
8 increments, while a 12-pitch line-feed
needs only 6.
.i0
.s3
Note that if your terminal has a plot switch, it must be turned
.it on
if
.it gsi
is to be used.

```

(where → means tab)

produces the following:

+12	permits use of 12-pitch, 6 lines/inch text. GSI terminals normally allow only two combinations: 10-pitch, 6 lines/inch, or 12-pitch, 8 lines/inch.
-n	controls the size of the half-line spacing. A half-line is, by default, equal to 4 vertical plot increments. A 10-pitch line-feed requires 8 increments, while a 12-pitch line-feed needs only 6.

Note that if your terminal has a plot switch, it must be turned *on* if *gsi* is to be used.

The effect of the `.lp` macro (which destroys the current tab settings as remembered by *nroff* or *troff*) must be turned off by inserting `.i0` before any line of input that is to be printed as regular text (i.e., indented 5 ens from the left margin and the subject headings). The `.i0` macro restores the tab settings to the default values (see below); it also causes a break.

To have the tag appear in bold when the manual page is *troffed*, the **.bn** command sequence:

```
.lp 14 7
.bn +12
permits use of 12-pitch . . .
```

The

would produce the argument "+12" in bold when *troffed*, and underscored when *nroffed*. Note the **.bn** macro *automatically* inserts the tab character after its argument.

BOLD OR ITALIC ITEMS

To make a word or phrase bold or italic in *troff*, the **.bd** and **.it** macros should be used. (Note that the *nroff* requests **.bd**, **.it**, and **.dt**—"embolden," "input-line count trap," and "diversion trap," respectively—are redefined here as **.Bd**, **.il**, and **.Dt**.) The *nroff* output for both the **.bd** and **.it** macros is underscored.

SETTING DEFAULT TABS

The **.dt** macro sets default tabs; it does *not* cause a break. Default tab stops are set every half-inch in *troff*; in *nroff*, they are set every eight characters.

SPECIAL CHARACTER STRINGS

Often it is necessary to use certain special characters when writing manual pages. To ensure consistency of usage for proper *nroff* and *troff* output, as well as for compatibility with existing manual pages, the following strings are defined:

<i>String</i>	<i>Output Character</i>
<code>*b</code>	beta (β)
<code>*a</code>	acute accent (´)
<code>*v</code>	bold vertical bar ()
<code>*p</code>	pi (π)
<code>*r</code>	registered symbol (®)
<code>*g</code>	grave accent (`)
<code>* </code>	1/6 em space in <i>troff</i> , one space in <i>nroff</i>
<code>*X</code>	multiply sign (\times)
<code>*u</code>	mu (μ)
<code>*></code>	right arrow (\rightarrow)
<code>*G</code>	capital gamma (Γ)

Note that to obtain the proper *troff* output for a minus sign (as opposed to a hyphen), it is necessary to type in the two-character sequence `\-` (see third line of Attachment 1).

OBTAINING OUTPUT

Manual pages may be printed on a terminal with the following command line:

```
nroff -ma files...
```

Nroff normally formats the output for a terminal in 10-pitch. To obtain *nroff* output in 12-pitch, the **-rV1** option must be added to the command line. Output to the Versatec printer is specified by the **-rV2** option as follows:

```
vp -b $bin\#$  "nroff -ma -rV2 files..."
```

where *bin#* is your output bin number. *Troff* output is obtained by:

```
troff -ma files...
```

EXAMPLES

Attachments 1 and 2 are examples of typical input and resulting output. Note that the *man(1)* command documented in the PWB/UNIX User's Manual (and given in the attachments) is to be used *only* to

alled manual pages, which are available on-line.

print the *X* *User's Manual—Edition 1.0*, T. A. Dolotta, R. C. Haight, and E. M. Piskorik, eds.

REF *Programmer's Manual, 6th edition* by K. Thompson and D. M. Ritchie.

ROFFITROFF User's Manual by J. F. Ossanna.

.th MAN I 1/10/77

.sh NAME

man \- print on-line documentation

.sh SYNOPSIS

.bd man

[options] documents

.sh DESCRIPTION

.it Man

is a shell command file that prints on-line documentation on the standard output by use of

.it nroff\c

(I) or

.it troff\c

(I).

On-line documentation consists of manual pages from the PWB/UNIX User's Manual.

.s3

The command line to print manual pages consists of:

.s3

.ti 10n

.bd man

[term]

[

.bd \-s

]

[section] title

.s3

where "term" is one of the following:

.s3

.lp 14 7

.bn \-t

produces photocomposed output;

.s3

.lp 14 7

.bn \-g

adapts the output to a DASI300 terminal in 12-pitch mode;

.s3

.lp 14 7

.bn \-450

adapts the output to a DASI450 terminal in 10-pitch mode;

.s3

.lp 14 7

.bn \-hp

adapts the output to a Hewlett-Packard terminal;

.s3

.lp 14 7

.bn \-v

followed by a space and a bin number, produces output on the Versatec printer.

Exactly one bin number

.it must

be specified when the

.bd \-v

option is used.

.i0

.s3

The

.bd \-s

option prints only the SYNOPSIS portion of a manual page.

.s3

.it Section

is the section number in the PWB/UNIX User's Manual in which the manual page is filed.

It is specified as a single Arabic decimal digit.

If

.it section

is omitted on the command line, the section number defaults to 1.

If the manual page is not in the given section, then a search is made in

.it all

sections of the manual.

If the page is not found (i.e., does not exist), an error message is produced.

.s3

.it Title

is the name of the manual page.

One or more titles may be specified in a single command.

.s3

Thus, the command line:

.s3

.ti 10n

man \-g l ed man tbl

.s3

would print out (in 12-pitch on a DASI300 terminal)

the currently installed versions of the manual pages for the commands

.it "ed, man,"

and

.it tbl,

all of which can currently be found in Section I of the manual.

.sh DIAGNOSTICS

.lp 35 30

``man page for xxx not found''→A manual page

(where → means tab)

for xxx does not exist.

.i0

.sh FILES

.lp 35 30

/usr/man/man[0-8]→Installed

PWB/UNIX documentation.

MAN (1)

PWB/UNIX 1/10/77

MAN (1)

NAME**man** – print on-line documentation**SYNOPSIS****man** [options] documents**DESCRIPTION**

Man is a shell command file that prints on-line documentation on the standard output by use of *nroff*(1) or *troff*(1). On-line documentation consists of manual pages from the PWB/UNIX User's Manual.

The command line to print manual pages consists of:

man [term] [-s] [section] title

where "term" is one of the following:

- t produces photocomposed output;
- g adapts the output to a DASI300 terminal in 12-pitch mode;
- 450 adapts the output to a DASI450 terminal in 10-pitch mode;
- hp adapts the output to a Hewlett-Packard terminal;
- v followed by a space and a bin number, produces output on the Versatec printer. Exactly one bin number *must* be specified when the -v option is used.

The -s option prints only the SYNOPSIS portion of a manual page.

Section is the section number in the PWB/UNIX User's Manual in which the manual page is filed. It is specified as a single Arabic decimal digit. If *section* is omitted on the command line, the section number defaults to 1. If the manual page is not in the given section, then a search is made in *all* sections of the manual. If the page is not found (i.e., does not exist), an error message is produced.

Title is the name of the manual page. One or more titles may be specified in a single command.

Thus, the command line:

man -g 1 ed man tbl

would print out (in 12-pitch on a DASI300 terminal) the currently installed versions of the manual pages for the commands *ed*, *man*, and *tbl*, all of which can currently be found in Section 1 of the manual.

DIAGNOSTICS

"man page for xxx not found" A manual page for xxx does not exist.

FILES

/usr/man/man[0-8]

Installed PWB/UNIX documentation.

