PLEXUS

Plexus Sys3 UNIX Programmer's Manual -- vol 2A

98-05036.5 Rev A   September 24, 1984

PLEXUS

PLEXUS

Sys3 UNIX   vol 2A   Programmer's Manual

Plexus Sys3 UNIX Programmer's Manual -- vol 2A

98-05036.5 Rev A   September 24, 1984

PLEXUS COMPUTERS, INC.

3833 North First St.

San Jose, CA 95134

408/943-9433

Programmer's Manual for UNIX* System III

Volume 2 - Supplementary Documents

January 1983

This volume contains documents that supplement the information contained in the *Plexus Sys3 UNIX Programmer's Manual - vol 1*. The documents are grouped roughly into the areas of General Works, Basics, Document Preparation Tools, Programming and Language Tools, and System Administration and Maintenance Tools. Further general information may be found in the July - August 1978 special issue of "The Bell System Technical Journal" on the UNIX Time Sharing System.

These documents contain occasional localisms, typically references to other operating systems like GCOS and IBM. In all cases, such references may be safely ignored by users of UNIX systems.

PREFACE

This manual contains a collection of documents that describe specific aspects of the UNIX* operating system. These include basic descriptions of the operating system, document preparation tools and programming and language tools.

Additional documents describing programming, language, administrative and maintenance tools are collected in the *Plexus Sys3 UNIX Programmer's Manual -- vol 2B* (Plexus publication number 98-05037).

Both these volumes (2A and 2B) should be used as supplementary documents for the *Plexus Sys3 UNIX Programmer's Manual -- vol 1A* (Plexus publication number 98-05045) and *Plexus Sys3 UNIX Programmer's Manual -- vol 1B* (Plexus publication number 98-05046). the basic reference manual for the operating system.

Comments
Please address all comments concerning this manual to:

Plexus Computers, Inc.
Technical Publications Dept.
2230 Martin Ave.
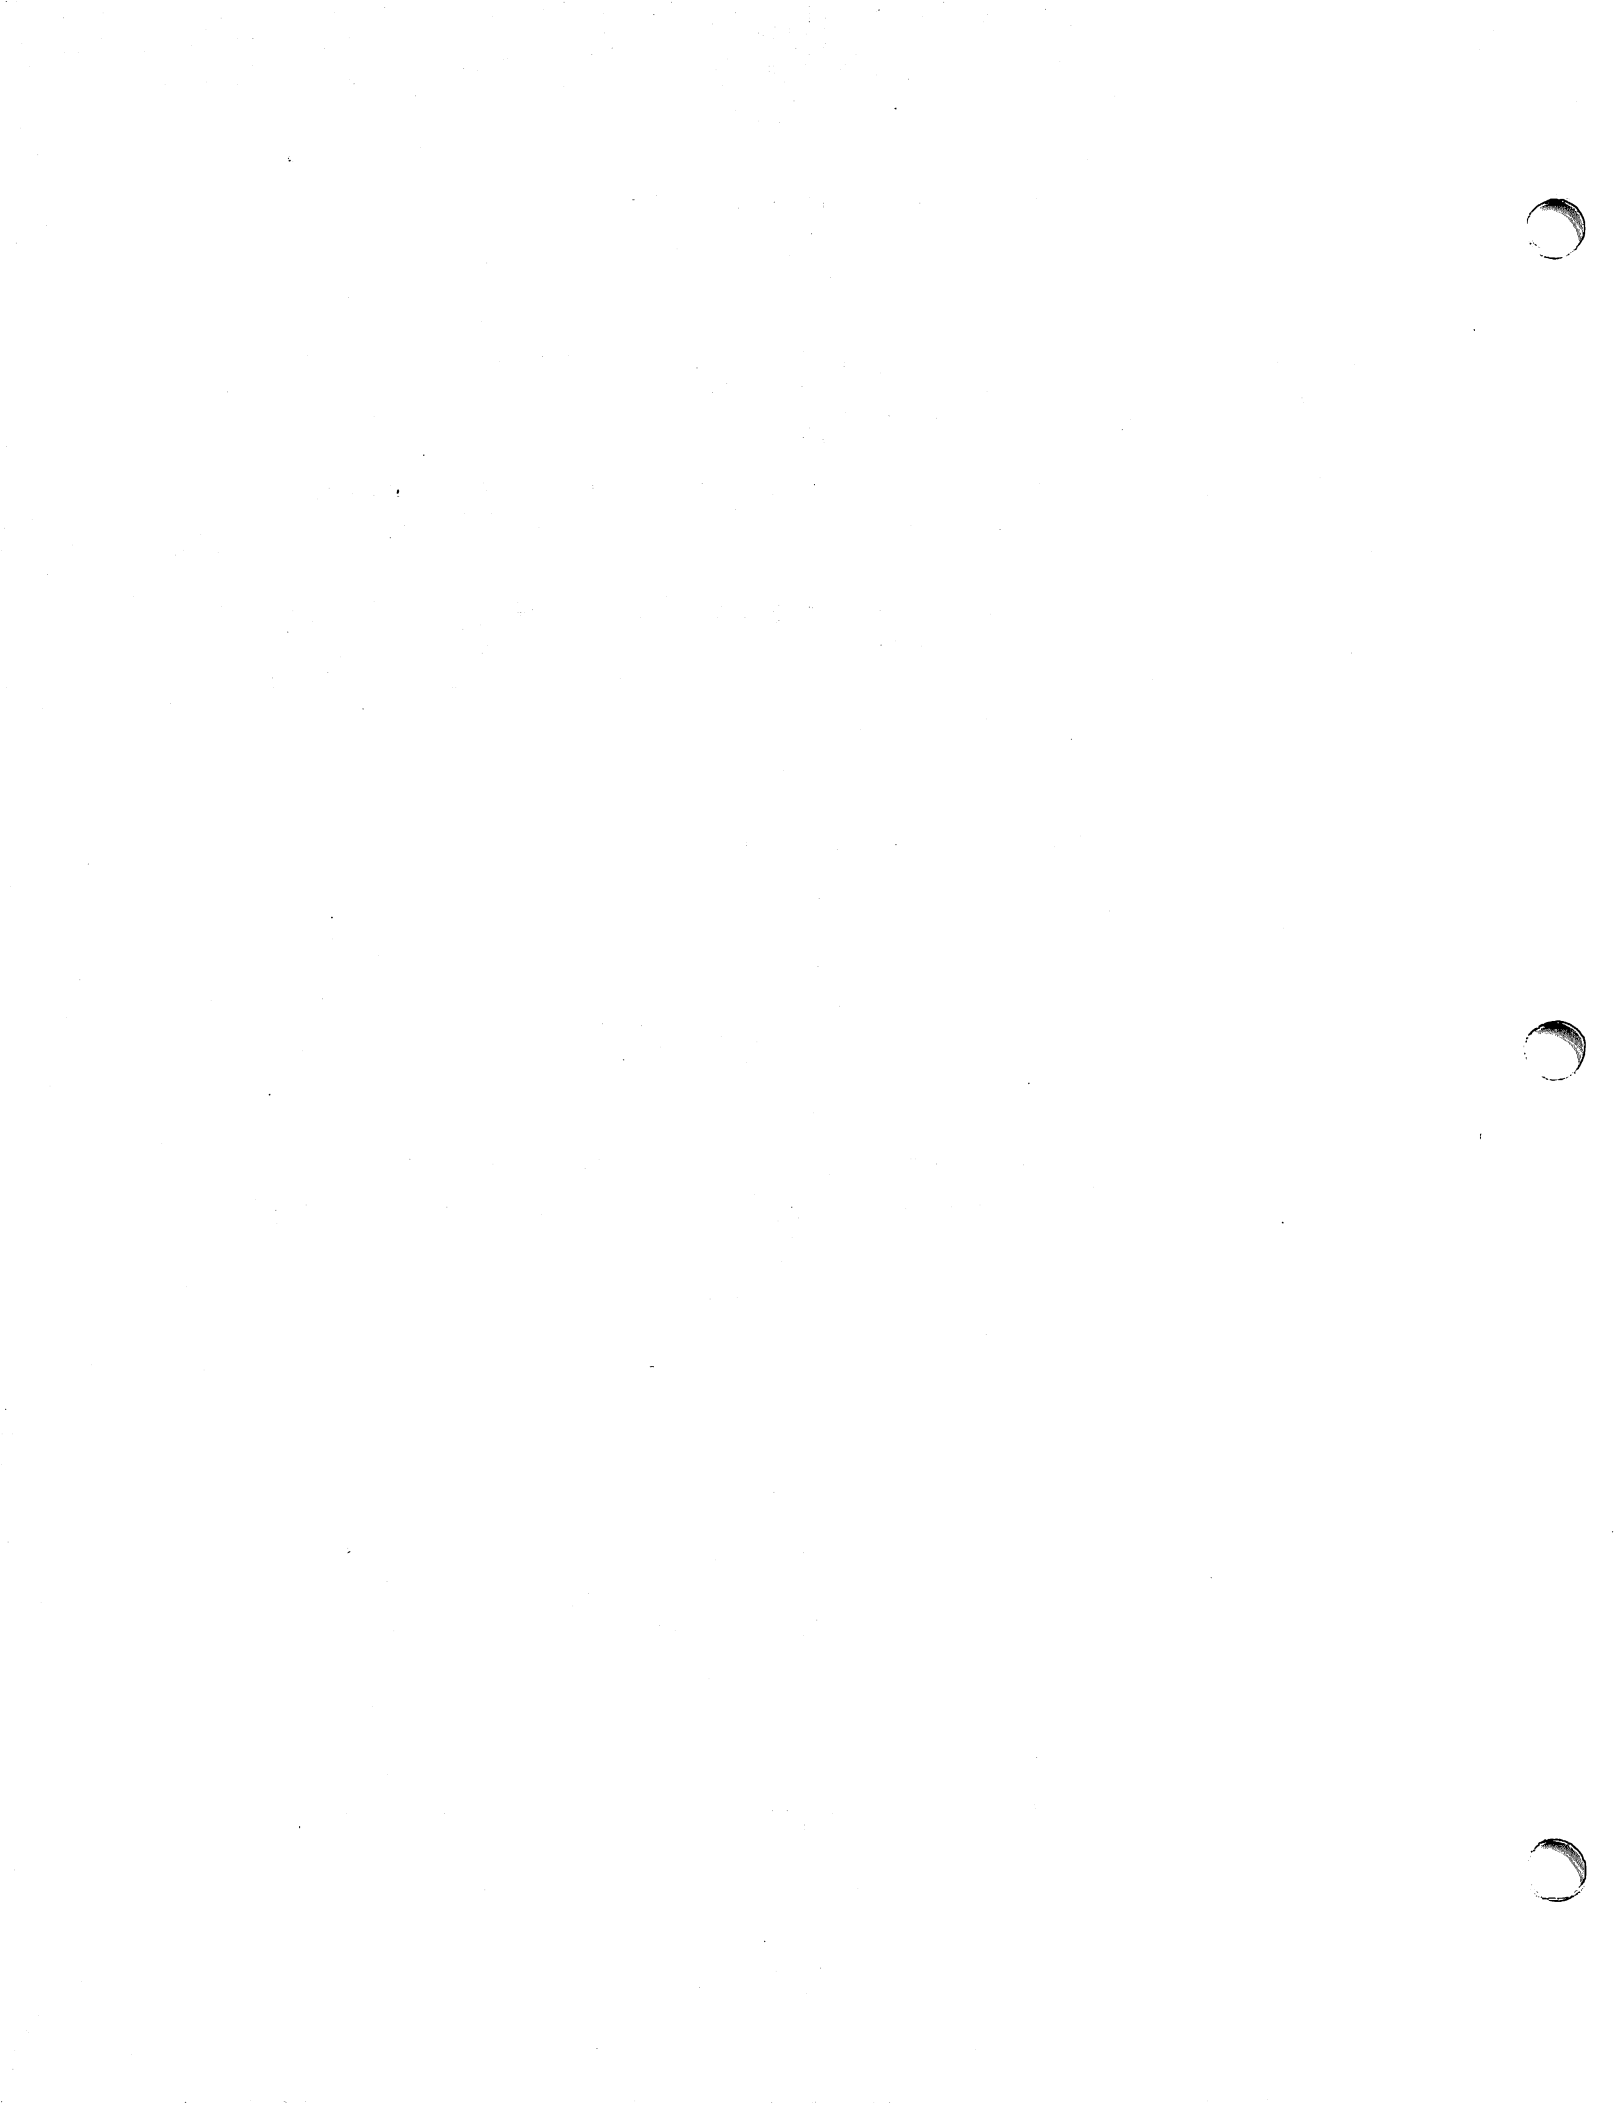Santa Clara, CA 95050
408/988-1755

Revision History
The second edition (#98-05036.2) contained a new Annotated Table of Contents.

The third edition (#98-05036.3) contained a new document, the *A68 Assembler Reference Manual*.

For this edition (#98-05036.4), several documents have been re-typeset.

ANNOTATED TABLE OF CONTENTS

General Works.

1. Overview

   - *Summary of UNIX System III.*

2. UNIX Time-Sharing System

   - *The UNIX Time-Sharing System.*
     D. M. Ritchie and K. Thompson.
     > An overview discussing the nature and implementation of the file system and of the user command interface.

Getting Started.

3. Roadmap

   - *UNIX/TS Documentation Road Map.*
     G. A. Snyder and J. R. Mashey.
     > An outline of important documents and information sources for new users.

4. UNIX for Beginners

   - *UNIX for Beginners - Second Edition.*
     B. W. Kernighan.
     > An introduction for the most basic use of the system.

5. Editor

   - *A Tutorial Introduction to the UNIX Text Editor.*
     B. W. Kernighan.
     > An easy way to get started with the text editor.

   - *Advanced Editing on UNIX.*
     B. W. Kernighan.
     > How secretaries, typists and programmers can make effective use of the UNIX facilities for preparing and editing text.

6. Shell

   - *An Introduction to the UNIX Shell.*
     S. R. Bourne.
     > An introduction to the use and capabilities of the command interpreter and programming language, the shell.

Document Preparation.

7. NROFF TROFF

- *NROFF/TROFF User's Manual.*
  J. F. Ossanna.

  A compact reference guide to the basic text formatting programs.

- *A TROFF Tutorial.*
  B. W. Kernighan.

  A beginner's guide to the use of TROFF for phototypsetting (and, by implication, NROFF).

8. Text Macro Packages

- *PWB/MM - Programmer's Workbench Memorandum Macros.*
  D. W. Smith and J. R. Mashey.

  User's guide and reference manual for PWB/MM, a general purpose package of text formatting macros for use with NROFF and TROFF.

9. TBL

- *Tbl - A Program to Format Tables.*

  Preprocessor for TROFF or NROFF that makes even very complex tables easy to specify.

10. EQN

- *Typesetting Mathematics - User's Guide (Second Edition).*
  B. W. Kernighan and L. L. Cherry.

  Describes the EQN and NEQN preprocessors for TROFF and NROFF, respectively. They allow inline typesetting complex formulae, equations, arrays, etc., displayed in a relatively simple command language.

- *A System for Typesetting Mathematics.*
  B. W. Kernighan and L. L. Cherry.

  Describes EQN, an easy-to-learn language for doing high-quality mathematical typesetting.

Programming.

11. UNIX Programming

- *UNIX Programming - Second Edition.*
  B. W. Kernighan and D. M. Ritchie.

  Describes how to write programs that interface to the operating system, either directly or through the standard I/O library.

12. C Language

- *The C Environment of UNIX TS.*
  A. R. Koenig.

  Describes the differences users may encounter when changing to UNIX TS from the various so-called "UNIX Sixth Edition" C compilers.

- *The C Programming Language - Reference Manual.*
  D. M. Ritchie.

  > Official statement of the syntax and semantics of C. Should be supplemented by "The C Programming Language", which contains a tutorial introduction and many examples. See Reference [1] below.

13. Lint

- *Lint, a C Program Checker.*
  S. C. Johnson.

  > Checks C programs for syntax errors, type violations, portability problems, and a variety of probable errors.

14. Make

- *Make - A Program for Maintaining Computer Programs.*
  S. I. Feldman.

  > Indispensable tool for making sure that alrge programs are properly compiled with minimum effort.

- *An Augmented Version of Make.*
  E. G. Bradford.

  > Describes an augmented version of the *make* command.

15. Debuggers

- *A Tutorial Introduction to ADB.*
  J. F. Maranzano and S. R. Bourne.

  > How to use ADB for debugging crashed systems and/or programs. Explains the various formatting options, techniques for debugging.

Supporting Tools and Languages.

16. Assembler

- *UNIX Z8000 Assembler Manual.*
  Craig C. Forney.

  > Describes the usage and input syntax of the UNIX Z8000 assembler (as).

17. AWK

- *Awk - A Pattern Scanning and Processing Language (Second Edition).*
  A. V. Aho, B. W. Kernighan, and P. J. Weinberger.

  > Makes it easy to specify many data transformation and selection operations.

18. Calculators

- *BC - An Arbitrary Precision Desk-Calculator Language.*
  L. L. Cherry and R. Morris.

  > A front end for DC (below) that provides infix notation, control flow, and built-in functions.

- *DC - An Interactive Desk Calculator.*
  L. L. Cherry and R. Morris.

  > Interactive desk calculator program that does arbitrary-precision integer arithmetic.

19. Fortran

- *A Portable Fortran 77 Compiler.*
  S. I. Feldman and P. J. Weinberger.

    Describes language compiled, interfaces between procedures, and file formats assumed by the I/O system. An appendix describes the Fortran 77 language.

- *RATFOR - A Preprocessor for a Rational Fortran.*
  B. W. Kernighan.

    Converts a Fortran with C-like control structures and cosmetics into real, ugly Fortran.

20. Graphics

- *PWB/Graphics Overview.*
  A. R. Feuer.

    Tells how to use a collection of numerical and graphical commands available to construct and edit numerical data plots and hierarchy charts. Best when with a Tektronix 4041 terminal.

- *Administrative Information for PWB/Graphics.*
  R. L. Chen and D. E. Pinkston.

    A reference guide for system administrators who are using or establishing a PWB/Graphics facility. It contains information about directory structure, installation, hardware requirements, etc., for PWB/Graphics.

- *A Tutorial Introduction to the Graphical Editor.*
  A. R. Feuer.

    How to make neat pictures and drawings with PWB Graphics and a Tektronix 4041 terminal.

21. LEX

- *Lex - A Lexical Analyzer Generator.*
  M. E. Lesk and E. Schmidt.

    Creates a recognizer for a set of regular expressions; each regular expression can be followed by arbitrary C code which is executed when the regular expression is found.

22. M4

- *The M4 Macro Processor.*
  B. W. Kernighan and D. M. Ritchie.

    A macro processor useful as a front end for C, Ratfor, Cobol, etc.

23. RJE

- *UNIX Remote Job Entry User's Guide.*
  A. L. Sabsevitz and K. A. Kelleman.

    The beginner's "how to" guide for sending a job to IBM or UNIVAC from UNIX via the Remote Job Entry facility.

- *UNIX Remote Job Entry Administrator's Guide.*
  M. J. Fitton.

    A system administrator's guide for setting up and managing a Remote Job Entry facility.

24. SED

- *SED - A Non-interactive Text Editor.*
  L. E. McMahon.

  > A variant of the text editor for processing large inputs and stream editing.

25. SCCS

- *Source Code Control System User's Guide.*
  L. E. Bonnani and C. A. Salemi.

  > Describes a collection of commands that controls retrieval of particular versions of files, administers changes to them, controls updating privileges to them, and records historical data about the changes.

- *Function and Use of an SCCS Interface Program.*
  L. E. Bonnani.

  > Describes the use of a Source Code Control System Interface Program to allow more than one user to use SCCS commands upon the same set of files.

26. UUCP

- *A Dial-Up Network of UNIX Systems.*
  D. A. Nowitz and M. E. Lesk.

  > Describes a design for implementing a dial-up network of systems for software transmission and distribution.

- *UUCP Implementation Description.*
  D. A. Nowitz.

  > Gives a detailed implementation description of UUCP for use by an administrator/installer of a UNIX system.

27. YACC

- *YACC: Yet Another Compiler-Compiler.*
  S. C. Johnson.

  > Converts a BNF specification of a language and semantic actions written in C into a compiler for the language.

28. VPM

- *Release 1.0 of the UNIX Virtual Protocol Machine.*
  P. F. Long and C. Mee, III.

  > Describes the initial release of the Virtual Protocol Machine, a new UNIX synchronous communication subsystem.

- *Release 2.0 of the UNIX Virtual Protocol Machine.*
  P. F. Long and C. Mee, III.

  > Describes the second release of the UNIX Virtual Protocol Machine.

Administration, Maintenance, and Implementation

29. Setting Up

   - *Setting Up UNIX.*
     R. C. Haight, T. J. Kowalski, M. J. Petrella, and L. A. Wehr.
     Procedures used to install UNIX System III and the steps necessary to regenerate the software.

30. Administrative Advice

   - *Administrative Advice for UNIX/TS.*
     R. C. Haight.
     Hints for getting UNIX up, getting it going, and keeping it going, plus some good stuff about hardware.

31. Accounting

   - *The PWB/UNIX Accounting System.*
     H. S. McCreary.
     Describes the structure, implementation, and management of the UNIX Accounting System.

32. FSCK

   - *FSCK - The UNIX/TS File System Check Program.*
     T. J. Kowalski.
     How to check and fix the file systems.

33. UNIX I/O

   - *The UNIX I/O System.*
     D. M. Ritchie.
     Provides guidance to writers of device driver routines, and is oriented more toward describing the environment and nature of device drivers.

34. UNIX Implementation

   - *UNIX Implementation.*
     D. M. Ritchie.
     How the system actually works inside.

35. C Compiler

   - *A Tour through the Portable C Compiler.*
     S. C. Johnson.
     How the portable C compiler works inside.

36. Security

   - *Password Security: A Case History.*
     R. Morris and K. Thompson.
     How the bad guys used to be able to break the password algorithm, and why they can't now, at least not so easily.
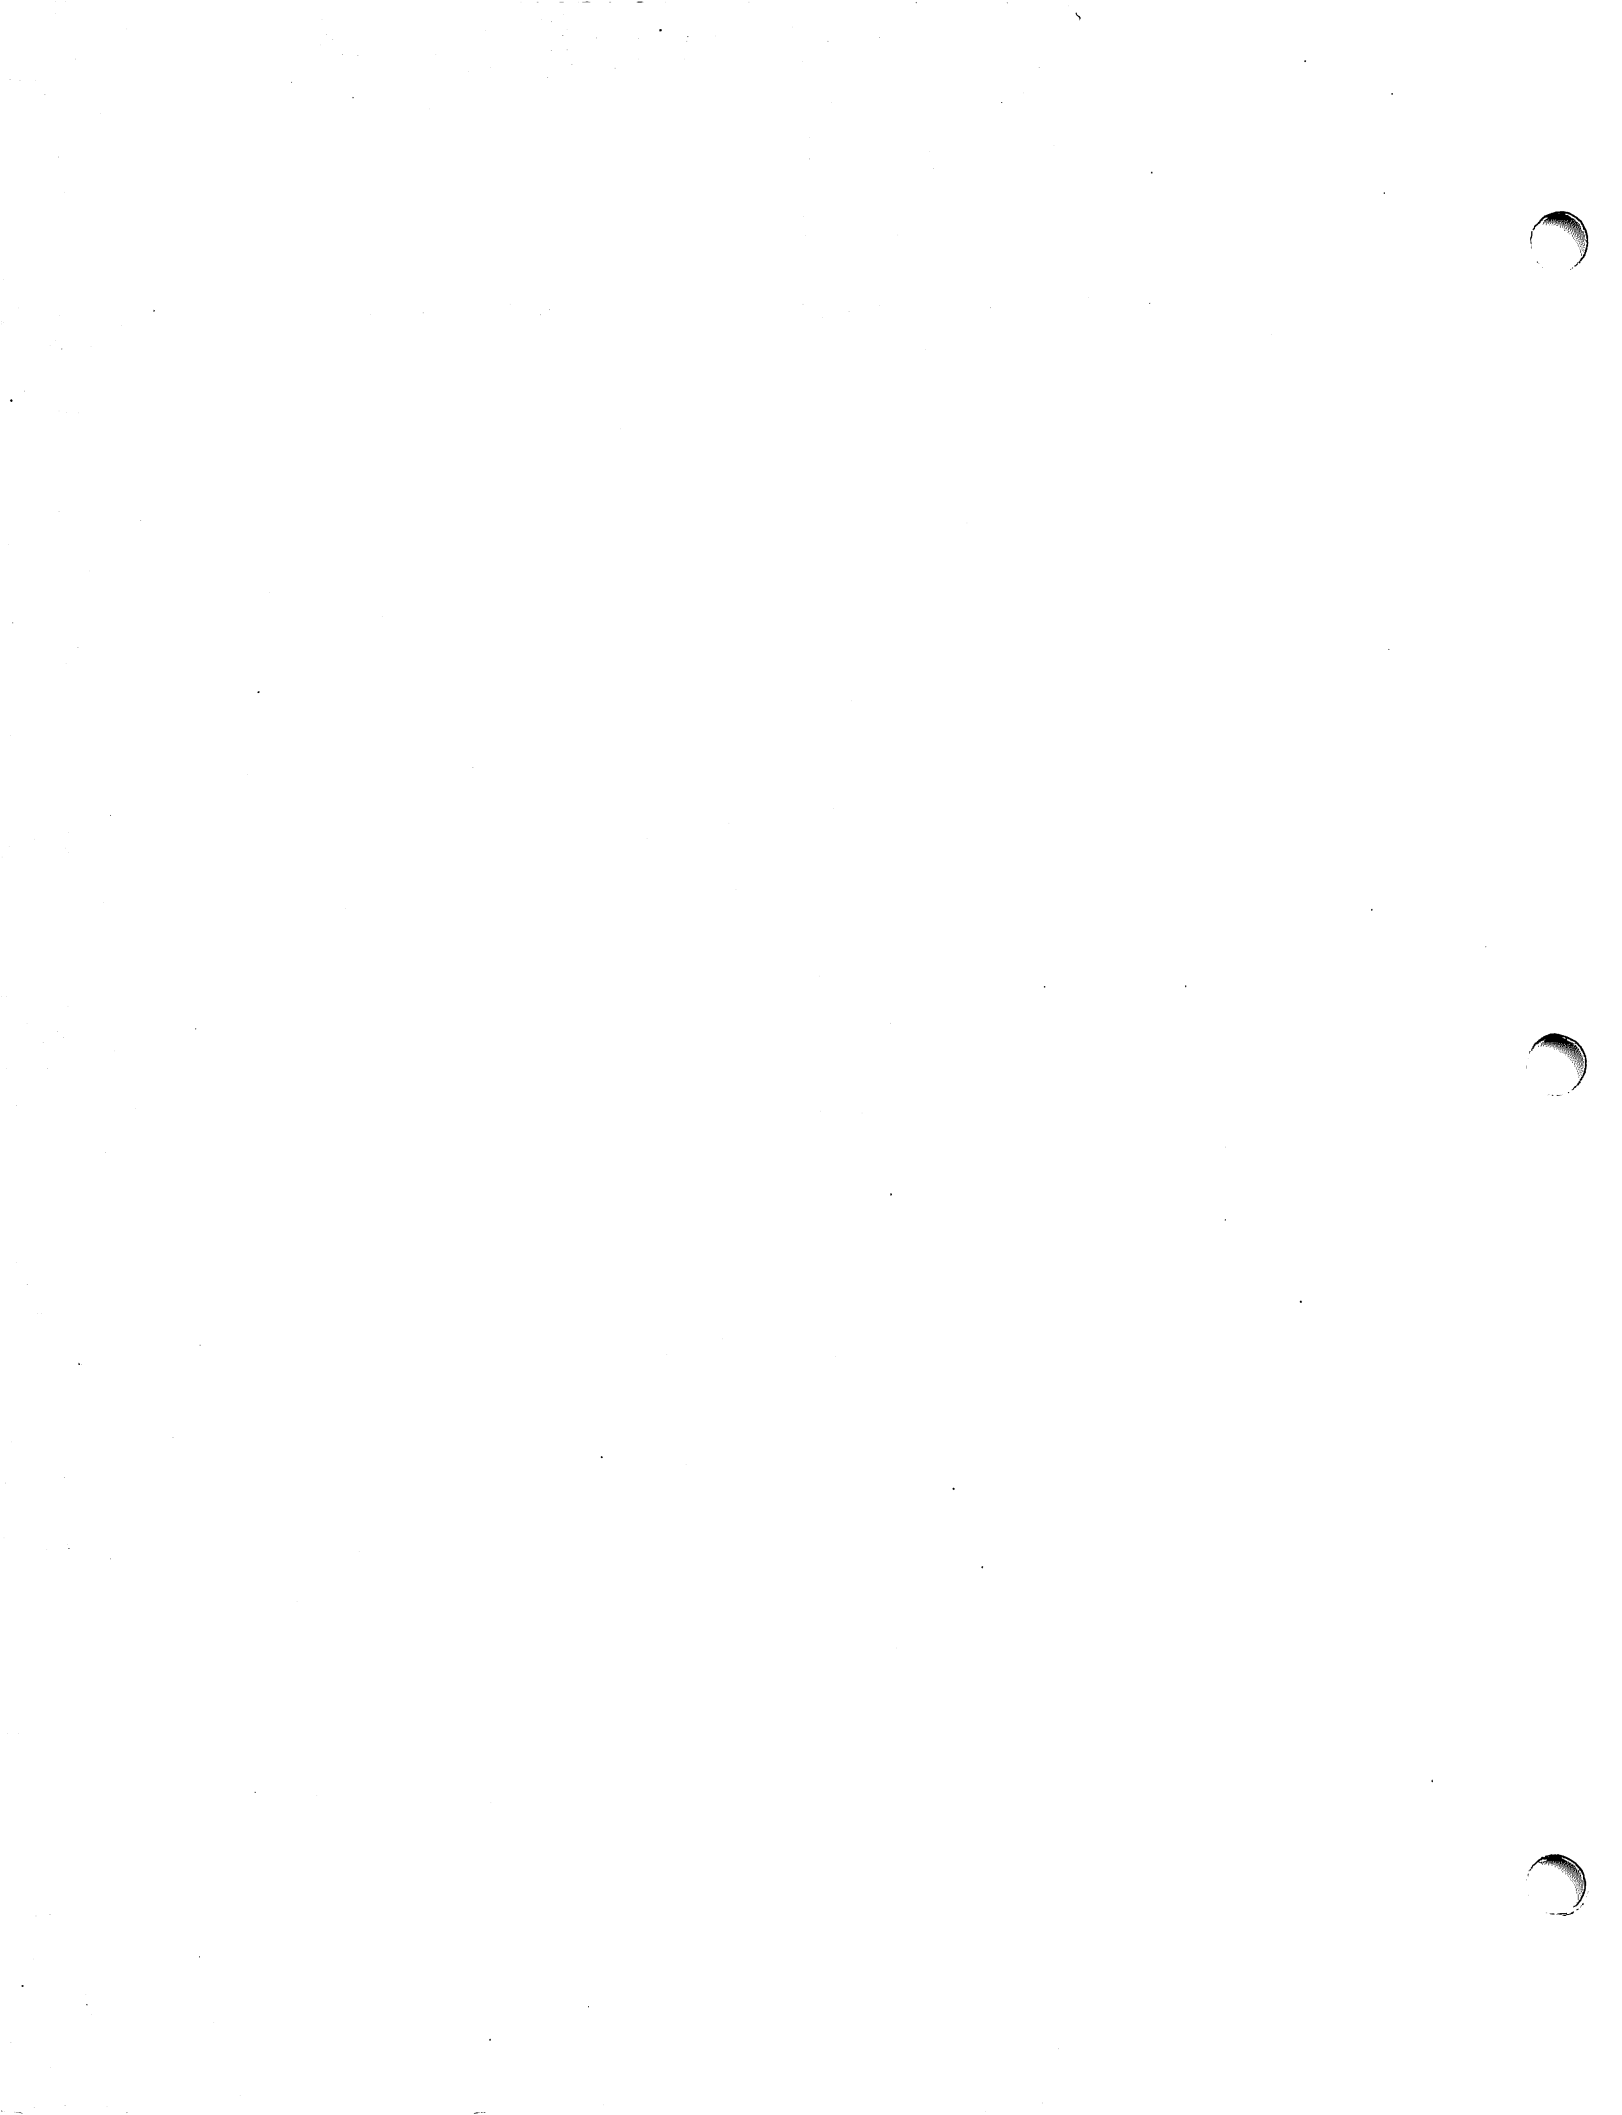
   - *On the Security of UNIX.*
     D. M. Ritchie.
     Hints on how to break UNIX, and how to avoid doing so.

References.

1. *The C Programming Language*, B. W. Kernighan and D. M. Ritchie, (Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1978).
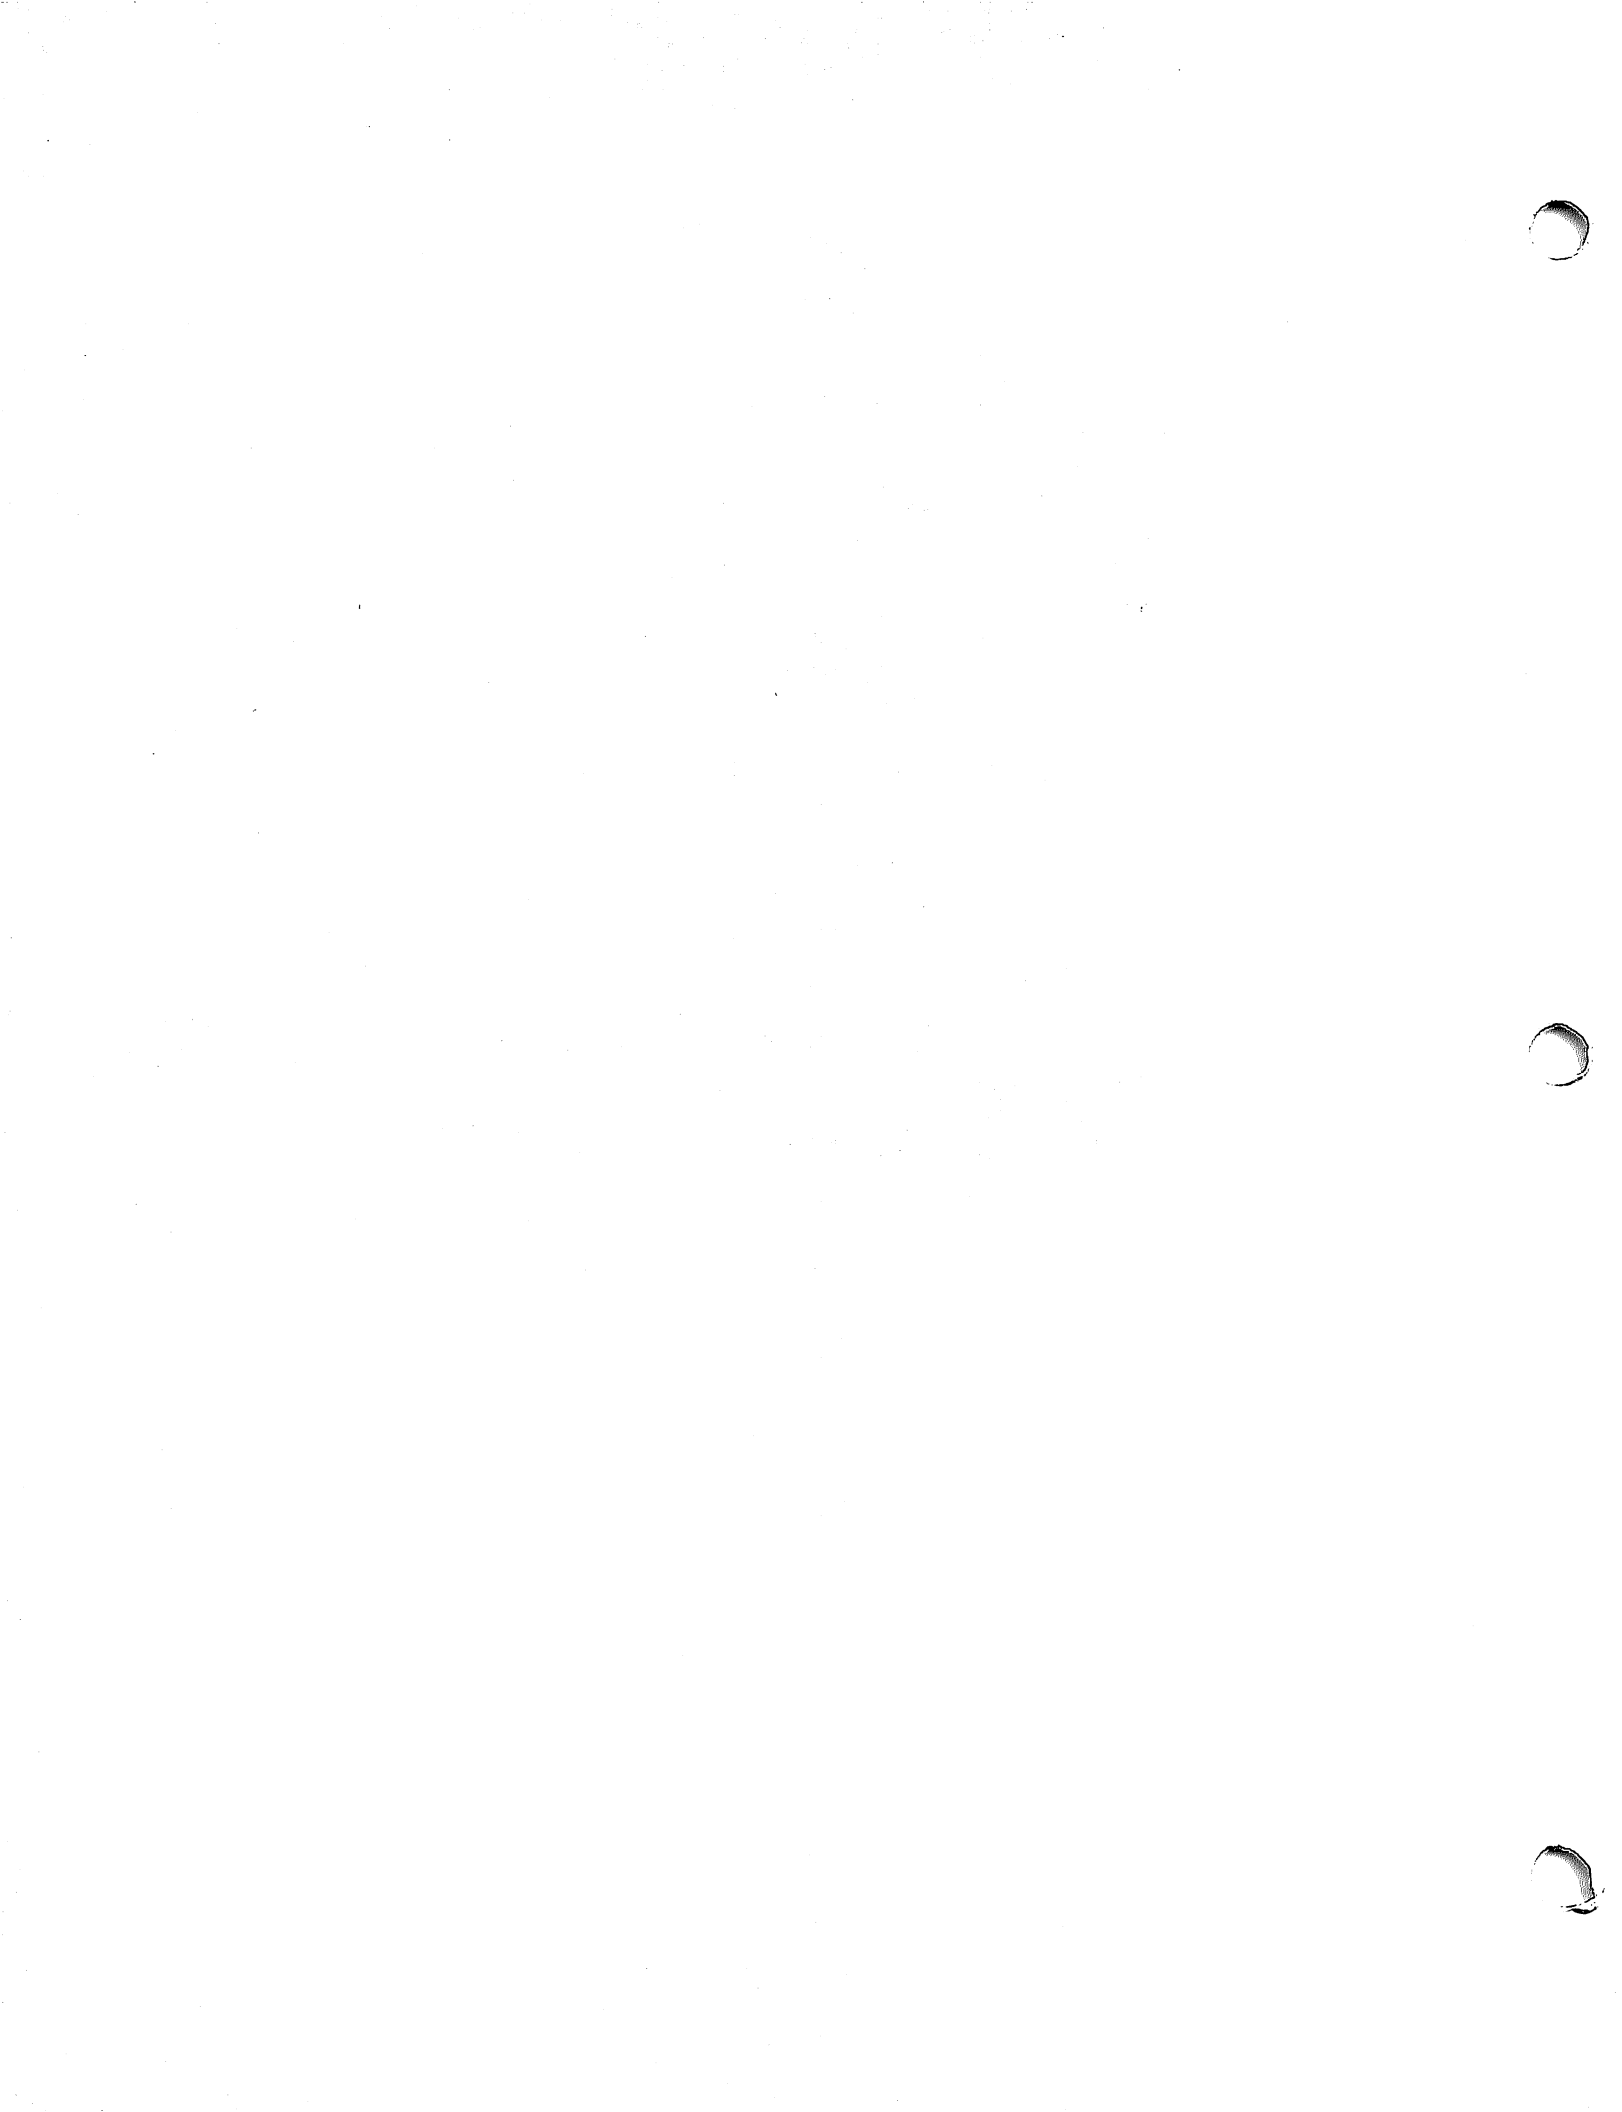
SUMMARY

OF

UNIX™ System III

CONTENTS

1. INTRODUCTION

UNIX* is a trademark for a family of computer time-sharing operating systems developed at Bell Laboratories and licensed by Western Electric. The first such operating system was conceived in 1969 when Bell Laboratories scientists set about to create a computing environment which was both comfortable to use and effective in their work. The result is an operating system of unusual simplicity, generality, and, above all intelligibility. The numerous software tools and utilities in UNIX Systems coupled with a simple user interface have propelled these systems to the forefront of popular operating systems in use both inside and outside the Bell System. Today, within the Bell System, UNIX Systems are used in applications involving office automation, electronic mail, document preparation, data access, data collection, networking, word processing and software development.

UNIX Systems are general-purpose, multi-user, interactive operating systems designed specifically to make the designer's, programmer's and documenter's computing environment simple, efficient, flexible and productive. Features of UNIX Systems include:

- Hierarchical, tree-structured file system

- Flexible, easy-to-use command language that can be tailored to meet specific user needs

- Ability to execute sequential, asynchronous, and background processes

- Powerful Text Editors

- High degree of portability

- Flexible, robust document preparation and text processing systems

- Access to the facilities of other ("host" or "target") computer systems, such as IBM MVS

- Support of numerous programming languages, including:

    - C - a high-level programming language conducive to structured programming

    - Fortran 77

    - SNOBOL

    - BS - An extension of Basic

    - Assembler

- Symbolic debugging systems

- Various system programming tools (e.g., compiler-compilers)

- Sophisticated "desk calculator" packages

- Graphics facilities

- Source Code Control System

- Network communications facilities between CPU's running under UNIX System III

- Synchronous Communications subsystem

- Resource Accounting for UNIX System III

- System Activity Package for monitoring CPU and disk/tape utilization, etc.

---

* UNIX is a trademark of Bell Laboratories

- Virtual Protocol Machine

UNIX System III, implemented throughout the Bell System in 1980, combines the features of UNIX Time-Sharing System, Seventh Edition and the PWB/UNIX system (Programmers Workbench) and incorporates numerous enhancements and new capabilities. The rapid growth in popularity of UNIX Systems in the office environment and the many new features provided by UNIX System III have prompted Western Electric to license this new release.

1.1 New Features and Enhancements

Major new features and enhancements provided by UNIX System III are:

- Software Support for the DEC** KMC11 Communications Microprocessor, permitting the downloading of many I/O functions from the main CPU:

  - KMC11/DZ11 Asynchronous Line Support, providing lower-overhead DMA processing for the DEC DZ11 asynchronous multiplexor.

  - A new multi-leaving IBM Remote Job Entry system with expanded utilities.

  - The Virtual Protocol Machine (VPM), providing a new synchronous communications interface for the implementation of link level protocols in a high level language.

- Numerous New Terminal Control Options

- Synchronous Terminal Interface for Teletype Model 40/4 Synchronous Terminals

- A new special file type called a FIFO (First-In/First-Out). FIFOs behave like pipes in UNIX Systems but have names and can be used to pass data between multiple, non-related processes, implement semaphores, etc.

- Parallel Communications Link driver, providing high performance local networking of up to 16 CPU's running under UNIX System III through the DEC PCL11-B.

- A New Systems Activity Package reports system statistics including CPU utilization, disk and tape I/O activities, terminal device activity, buffer usage, system calls, etc.

- Enhanced C Compiler, Text Processing software, Source Code Control System and Graphics capabilities.

Additionally, numerous enhancements to existing software tools, new tools, games, etc. are included.

** DEC is a trademark of Digital Equipment Corporation

2. PRODUCT OVERVIEW OF UNIX SYSTEM III

A synopsis of the major features of UNIX System III follows.

2.1 The File System

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

- File systems are addressable to one billion bytes

- Simple and consistent naming conventions; name 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

- 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, and for all other users; protection modes can be set dynamically.

- Facilities for creating, accessing, moving and processing files, directories or sets of these in a simple, uniform and natural way.

- Provides device-independence since each physical I/O device, from interactive terminals to main memory, is treated like a file, allowing uniform file and device I/O.

2.2 The Command Language

Communication with a CPU running under UNIX System III is normally carried out with the aid of a program called "shell". The shell is both a command language interpreter and a programming language that provides an interface to the operating system. As a command language interpreter it reads lines typed by the user and interprets them as requests to execute other programs.

As a programming language the shell includes such features as:

- Control-flow primitives

- Parameter Passing

- Variables and string substitution

- Constructs such as "while", "if then else", "case" and "for"

- Two-way communication between shell commands

The shell can modify the environment in which commands run. Input and output can be redirected to files and processes that communicate through "pipes" and semaphores can be invoked. Commands are found by searching directories in the file system in a sequence that can be defined by the user. Command lines can be read either from the terminal or from a file, thus allowing command procedures to be stored for later use.

By utilizing the shell as a programming language, users of UNIX System III can eliminate a great deal of the programming drudgery that often accompanies a project. Many manual procedures can be quickly, inexpensively and conveniently automated. Because it is so easy to create and use shell procedures, projects using UNIX System III can customize the general environment into one tailored to its own requirements, organizational structure and terminology.

2.3 Document Preparation and Text Processing

UNIX systems are renowned for their text processing and document preparation facilities. Included are a powerful text editor, programmable text formatters that are robust and yet easy to use, text processing macro packages which simplify the creation of letters, reports, memoranda, books, etc., special processors for mathematical expressions and tabular material, and numerous supporting utilities such as spelling and typographical error detectors. A variety of output devices are supported, including phototypesetters, typewriter-like terminals, line printers and dot-matrix printer-plotters. UNIX System III contains more than one hundred enhancements to these facilities, improving their efficiency, flexibility and robustness.

2.4 Remote Job Entry

The Remote Job Entry (RJE) facility in UNIX System III provides for the submission and retrieval of jobs by a remote computer running under UNIX System III from an IBM host system. To the host system, such a remote system appears as an IBM System/360 remote multi-leaving work station. At the request of a user of the remote system, RJE gathers the jobstream to be sent to the host system and subsequently retrieves from the host the resulting output, notifies the user of the output's arrival and places it in a convenient file on the remote system for perusal or optionally passes it as input to a program in UNIX System III to be automatically executed. Additionally, an interface utility is provided to execute JES or HASP commands on the host computer from the remote system.

This release contains a new, more efficient KMC-based, multi-leaving RJE system and also provides a DQS-base RJE for upward compatibility.

2.5 Source Code Control System

The Source Code Control System (SCCS) in UNIX System III is an integrated set of commands designed to aid software development projects or document preparation groups *control* changes to source code or files of text (e.g. manuals). SCCS provides facilities for storing, updating, and retrieving, by version number or date, all versions of source code modules or of a document, and for recording who made the change, when it was made and why. SCCS is designed to solve most of the source code and documentation control problems that 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 versions of a source module or document are stored together, so that common code or text is stored only once.

- Versions of a source module or document can be protected from unauthorized changes.

- Sufficient identifying information (module name, date, time, etc.) can be automatically inserted into the source module or document, enabling one to identify the exact version of a module given only the corresponding load module, memory dump or text.

UNIX System III includes several enhancements and new features for SCCS.

2.6 The Virtual Protocol Machine

The Virtual Protocol Machine (VPM) in UNIX System III is a general-purpose synchronous communications interface that allows link-level protocols to be implemented in a high-level language. The VPM software consists of a protocol compiler, a driver, an interpreter that executes in the KMC11-B microprocessor, and several utility programs. Byte-oriented, half-duplex synchronous (BISYNC) and bit-oriented, full-duplex protocols (e.g. HDLC, X.25) can be implemented.

2.7 File transfer between CPU's under UNIX System III (UUCP)

UUCP consists of a series of programs designed to permit communications between CPU's running under UNIX System III using either dial-up or hard-wired communication lines. UUCP's simple user interface provides for file transfers and remote command execution. UUCP operates in batch mode; files are created in a spool directory for subsequent processing by a UUCP daemon program, thus eliminating the need for tying up a user port while the actual file transfer takes place.

2.8 Resource Accounting in UNIX System III

The Resource Accounting System in UNIX System III collects per-process resource utilization data, records user connect time, monitors disk utilization and charges fees to individual users. Report generators are provided for data reduction of the accounting data and generation of summary files and reports.

2.9 The System Activity Package

A number of counters have been incorporated into UNIX System III that record various system activities, including CPU utilization, disk and tape I/O, buffer cache hit ratios, etc. Report software depict this data in tabular and graphic formats. The System Activity Package is particularly useful for system administrators to monitor performance of a CPU running under UNIX System III, thus enabling them to configure their system for optimum throughput.

2.10 Graphics Facilities

UNIX System III includes non-interactive graphics facilities consisting of commands that construct data plots, pie charts, bar charts, histograms, etc. An interactive graphical editor is provided to construct, edit and display graphic images.

3. COMPUTER CONFIGURATIONS FOR UNIX SYSTEM III

UNIX System III is available for the following DEC processors:

> PDP 11/23
> PDP 11/34
> PDP 11/44
> PDP 11/45
> PDP 11/70
> VAX 11/780

This wide range of processing price/performance coupled with the ease of porting software developed on one computer to other computers with UNIX System III provides enormous flexibility in planning a computing environment. UNIX System III can be tailored to meet the needs of user communities ranging in size from one user up to 48 simultaneous users on a single system. The cost per user hour of UNIX System III is significantly lower than that of most other interactive computer systems. The processors and peripherals supported by UNIX System III are depicted in Table 1. Each hardware configuration must include at least the following:

- A processor with memory management
- Memory as specified in Table 1
- A system disk
- A magnetic tape drive (Except for the PDP 11/23)
- System clock
- System console

The following equipment is strongly recommended:
- Floating-point hardware
- Communications controller
- Full-duplex 96 character ASCII terminals
- Extra disk drive for system backup
- Synchronous communication interface for RJE

TABLE 1

Processors and Peripherals Supported by UNIX System III

Processor	PDP-11/23	PDP-11/34	PDP-11/45	PDP-11/44	PDP-11/70	VAX-11/780
Memory	256KB	256KB	256KB	.05-1MB	0.5-1MB	1-2MB
Floating Point:	—	FP11-A	FP11-B	—	FP11-C	FP780
Console Terminal:						
LA36	S	S	S	S	S	S
LA120	F	F	F	F	F	F
Disk drives:						
RL01/2 (2 required)	F	F	S	S	N	—
RK05 (2 required)	—	O	O	O	—	—
RP03	—	—	O	O	—	—
RP04,05	—	O	O	—	O	—
		(UNIBUS)	(UNIBUS)	(UNIBUS)	(MASSBUS)	
RP06	—	S	F	—	F	F
		(UNIBUS)	(UNIBUS)	(UNIBUS)		(MASSBUS)
Fixed-head disks:						
RF11	—	—	O	—	—	—
RS03,4	—	—	N	N	N	—
			(UNIBUS)	(UNIBUS)	(MASSBUS)	
Tape drives:						
TU10	—	O	O	O	—	—
TU16	—	O	O	O	O	—
TE16	—	F	F	F	F	F
TU45	—	—	—	—	S	S
Line printer:						
LP11 (or equiv.)†	—	S	S	S	S	S
Asynch. interfaces:						
DLV11	S	—	—	—	—	
DL11-E	—	S	S	S	S	—
DZ11	—	S	S	S	S	S
KMC11B/DZ11‡	—	F	F	F	F	F
DH11	—	O	O	O	O	—
Synch. interfaces:						
KMC11B/DMC	—	F	F	F	F	F
DMC11	—	S	S	S	S	S
DU11	—	O	O	O	O	—
DQS11B	—	—	O	O	O	—
Auto Call Unit:						
DN11AA/DA	—	S	S	S	S	S
Parallel interface:						
DRV11	S	—	—	—	—	
DR11C	—	S	S	S	S	—
DR11B	—	S	S	S	S	—
DA11B (use DMCs)	—	N	N	N	N	—
PCL11B	—	N	N	N	N	—
Max. Users	4	8	16-24	20-30	32-40	48+

Key: F — First choice.
      S — Supported.
      O — Driver provided, but device is obsolete.
      N — Driver provided, but device is not recommended.
      — — No driver - not supported.
† Use a printer on an RS232 interface.
‡ 4 DZ11s (32 lines) per KMC.

October 1981

4. SOFTWARE

Most of the programs available as commands in UNIX Systems are listed. Source code and printed manuals are distributed for all of the listed software except games. Almost all of the code is written in C. Commands are self-contained and do not require extra setup information, unless specifically noted as "interactive." Interactive programs can be made to run from a prepared script simply by redirecting input. Most programs intended for interactive use (e.g., the editor) allow for an escape to command level (the Shell). Most file processing commands can also go from standard input to standard output ("filters"). The piping facility of the Shell may be used to connect such filters directly to the input or output of other programs.

4.1 Basic Software

This includes the time-sharing operating system with utilities, a machine language assembler and a compiler for the programming language C - enough software to write and run new applications and to maintain or modify UNIX System III itself.

4.1.1 Operating System

UNIX   The basic resident code on which everything else depends. Supports the system calls, and maintains the file system. A general description of the design philosophy and system facilities in UNIX Systems appeared in the Communications of the ACM, July, 1974. A more extensive survey is in the Bell System Technical Journal for July-August 1978. Capabilities include:

- Reentrant code for user processes.

- Separate instruction and data spaces.

- "Group" access permissions for cooperative projects, with overlapping memberships.

- Alarm-clock timeouts.

- Timer-interrupt sampling and interprocess monitoring for debugging and measurement.

- Multiplexed I/O for machine-to-machine communication.

DEVICES   All I/O is logically synchronous. I/O devices are simply files in the file system. Normally, invisible buffering makes all physical record structure and device characteristics transparent and exploits the hardware's ability to do overlapped I/O. Unbuffered physical record I/O is available for unusual applications. Drivers for these devices are available; others can be easily written:

- Asynchronous interfaces: DZ11, KMC11B/DZ11, DH11, DL11.

- Support for most common ASCII terminals.

- Synchronous interface: KMC11B/DMC11, DMC11, DQS11, DU11.

- Automatic calling unit interface: DN11.

- Parallel interface: DR11C/B, PCL11B.

- Line printer: LP11, or equivalent.

- Magnetic tape: TU10, TU16, TU45, TU77, or equivalent.

- DECtape: TC11.

- Fixed head disk: RS11, RS03 and RS04, or equivalent.

- Pack type disk: RP03, RP04, RP05, RP06, or equivalent; minimum-latency seek scheduling.

- Cartridge-type disk: RL01, RK05, one or more physical devices per logical device.

- Null device.

- Physical memory of PDP-11, or mapped memory in resident system.

- Phototypesetter: Graphic Systems System/1 through DR11C.

BOOT   Procedures to get UNIX System III started.

CONFIG   Tailor device-dependent system code to hardware configuration. As distributed, UNIX System III can be brought up directly on any acceptable CPU with any acceptable disk, any sufficient amount of core, and either clock. Other changes, such as optimal assignment of directories to devices, inclusion of floating point simulator, or installation of device names in file system, can then be made at leisure.

4.1.2 User Access Control

LOGIN   Sign on as a new user.

- Verify password and establish user's individual and group (project) identity.

- Adapt to characteristics of terminal.

- Establish working directory.

- Announce presence of mail (from MAIL).

- Publish message of the day.

- Execute user-specified profile.

- Start command interpreter or other initial program.

PASSWD   Change a password.

- User can change his own password.

- Passwords are kept encrypted for security.

NEWGRP   Change working group (project). Protects against unauthorized changes to projects.

4.1.3 Terminal Handling

TAB   Set tab stops appropriately for specified terminal type.

STTY   Set 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.

- Half vs. full duplex.

- Carriage return + line feed vs. newline.

- Interpretation of tabs.

- Parity.

- Mapping of upper case to lower.

- Raw vs. edited input.
- Delays for tabs, newlines and carriage returns.

4.1.4 File Manipulation

CAT	Concatenate one or more files onto standard output. Particularly used for unadorned printing, for inserting data into a pipeline, and for buffering output that comes in dribs and drabs. Works on any file regardless of contents.
CP	Copy one file to another, or a set of files to a directory. Works on any file regardless of contents.
PR	Print files with title, date, and page number on every page.

- Multicolumn output.
- Parallel column merge of several files.

LPR	Off-line print. Spools arbitrary files to the line printer.
CMP	Compare two files and report if different.
TAIL	Print last n lines of input.

- May print last n characters, or from n lines or characters to end.

CSPLIT

SPLIT	Split a large file into more manageable pieces. Occasionally necessary for editing (ED).
DD	Physical file format translator, for exchanging data with foreign systems, especially IBM 370's.
SUM	Sum the words of a file.
PACK	Store files in a compressed form.
UNPACK	Restore compressed files to original form.
PCAT	Does for packed files what CAT does for ordinary files.
CPIO	Copy file archives in and out.
PASTE	Merge corresponding lines of several files or subsequent lines of a single file.
CUT	Cut out selected fields of each line of a file.
REFORM	Reformat text files.

- Rearrangment of tab characters.
- Trim trailing blanks.
- Truncate lines.
- Prepend blank lines.

TOUCH	Update access and modification times of a file.

4.1.5 Manipulation of Directories and File Names

RM	Remove a file. Only the name goes away if any other names are linked to the file.

- Step through a directory deleting files interactively.
- Delete entire directory hierarchies.

LN   "Link" another name (alias) to an existing file.

LINK

UNLINK   Exercise link and unlink system calls.

MV   Move a file or files. Used for renaming files.

CHMOD   Change permissions on one or more files. Executable by files' owner.

CHOWN   Change owner of one or more files.

CHGRP   Change group (project) to which a file belongs.

CHROOT   Change root directory for a command.

MKDIR   Make a new directory.

RMDIR   Remove a directory.

MVDIR   Move a directory.

CD   Change working directory.

UMASK   Set file-creation mode mask.

FIND   Prowl the directory hierarchy finding every file that meets specified criteria.

- Criteria include:

  name matches a given pattern,

  creation date in given range,

  date of last use in given range,

  given permissions,

  given owner,

  given special file characteristics,

  boolean combinations of above.

- Any directory may be considered to be the root.
- Perform specified command on each file found.

4.1.6 Running of Programs

SH   The Shell, or command language interpreter.

- Supply arguments to and run any executable program.
- Redirect standard input, standard output, and standard error files.
- Pipes: simultaneous execution with output of one process connected to the input of another.
- Compose compound commands using:

  if . . .then . . . else,

  case switches,

while loops,

for loops over lists,

break, continue and exit,

parentheses for grouping.

- Initiate backgound processes.

- Perform Shell programs, i.e., command scripts with substitutable arguments.

- Construct argument lists from all file names satisfying specified patterns.

- Take special action on traps and interrupts.

- User-settable search path for finding commands.

- Executes user-settable profile upon login.

- Optionally announces presence of mail as it arrives.

- Provides variables and parameters with default setting.

RSH — Restricted version of the shell. Allows limited command execution.

ENV — Set an environment for command execution.

TEST — Tests for use in Shell conditionals.

- String comparison.

- File nature and accessibility.

- Boolean combinations of the above.

EXPR — String computations for calculating command arguments.

- Integer arithmetic.

- Pattern matching.

WAIT — Wait for termination of asynchronously running processes.

LINE — Read a line from terminal, for interactive Shell procedure.

ECHO — Print remainder of command line. Useful for diagnostics or prompts in Shell programs, or for inserting data into a pipeline.

SLEEP — Suspend execution for a specified time.

NOHUP — Run a command immune to hanging up the terminal.

NICE — Run a command in low (or high) priority.

KILL — Terminate named processes.

CRON — Schedule regular actions at specified times.

- Actions are arbitrary programs.

- Times are conjunctions of month, day of month, day of week, hour and minute. Ranges are specifiable for each.

TEE — Pass data between processes and divert a copy into one or more files.

4.1.7 Status Inquiries

LS	List the names of one, several, or all files in one or more directories.

- Alphabetic or temporal sorting, up or down.
- Optional information: size, owner, group, date last modified, date last accessed, permissions, i-node number.

FILE	Try to determine what kind of information is in a file by consulting the file system index and by reading the file itself.
DATE	Print today's date and time. Has considerable knowledge of calendric and horological peculiarities.

- May set date and time.

DF	Report amount of free space on file system devices.
DU	Print a summary of total space occupied by all files in a hierarchy.
PS	Report on active processes.

- List your own or everybody's processes.
- Tell what commands are being executed.
- Optional status information: state and scheduling info, priority, attached terminal, what it's waiting for, size.

WHO	Tell who's on the system.

- List of presently logged users, ports and times on.
- Optional history of all logins and logouts.

WHODO	Tell who's doing what on the system. Prints merged, reformatted, and dated output from the WHO and PS commands.
ID	Print user and group IDs and names.
IOSTAT	Print statistics about system I/O activity.
TTY	Print name of your terminal.
PWD	Print name of your working directory.
LOGNAME	Print login name.
UNAME	Print the current system name of UNIX System III.
SYSDEF	Extract configuration information from the operating system file.
DEVNAM	Identify the special file associated with a mounted file system.

4.1.8 Backup and Maintenance

MOUNT	Attach a device containing a file system to the tree of directories. Protects against nonsense arrangements.
UMOUNT	Remove the file system contained on a device from the tree of directories. Protects against removing a busy device.
SETMNT	Create the "mnttab" table used by the MOUNT and UNMOUNT commands.
MKFS	Make a new file system on a device.

MKNOD	Make an i-node (file system entry) for a special file. Special files are physical devices, virtual devices, physical memory, etc.
TP	
TAR	Manage file archives on magnetic tape of DECtape. TAR is newer.

- Collect files into an archive.
- Update DECtape archive by date.
- Replace or delete DECtape files.
- Print table of contents.
- Retrieve from archive.

DUMP	Dump the file system stored on a specified device, selectively by date, or indiscriminately.
RESTOR	Restore a dumped file system, or selectively retrieve parts thereof.
SU	Temporarily become the super user with all the rights and privileges thereof. Requires a password.
FSCK	Audit and interactively repair inconsistent conditions for file systems.

- For consistent file systems, print gross statistics; number of files, number of blocks used, number of blocks free.
- For inconsistent file systems, the operator is prompted for concurrence before corrections are attempted.

NCHECK	Generate file names from i-numbers.
CLRI	Peremptorily expunge a file and its space from a file system. Used to repair damaged file systems.
SYNC	Force all outstanding I/O on the system to completion. Used to shut down gracefully.
FSDB	File system debugger. Used to patch up a damaged file system after a crash.
VOLCOPY	Copy a file system with label checking.
LABELIT	Create initial labels for disk or tape file systems.
INSTALL	Install commands.
SHUTDOWN	Terminate all processing.

- Broadcast log off message.
- Unmount file systems.
- Run file-save procedure.
- Update super blocks.
- Stop processing.

4.1.9 Accounting and System Activity Reporting

The accounting software is structured as a set of tools (consisting of both C programs and shell procedures) that can be used to build accounting systems.

ACCTON	Turn process accounting on (or off).
ACCTDISK	Build accounting records containing user ID, login name, and number of disk blocks.
ACCTDUSG	Compute disk resource consumption by login.
ACCTWTMP	Write an accounting record.
ACCTCMS	Print command summary from per-process accounting records.
ACCTCON	Connect-time accounting. Create accounting records for login/logoff activity.
ACCTMERG	Merge or add total accounting records.
ACCTPRC	Process accounting. Creates accounting records for executed processes.
ACCTSH	A collection of shell procedures to manage accounting activities.
RUNACCT	Run daily accounting. The main daily accounting shell procedure. Normally initiated via CRON.
FWTMP	
WTMPFIX	Commands for manipulating accounting file records.
PROFILER	A system of commands to facilitate an activity study of UNIX System III.
SAG	Display, in graphical form, system activity of UNIX System III for a specified time interval.
TIMEX	Time a command and generate a system activity report.

4.1.10 Communication

MAIL	Mail a message to one or more users. Also used to read and dispose of incoming mail. The presence of mail is announced by LOGIN and optionally by SH.
	• Each message can be disposed of individually.
	• Messages can be saved in files or forwarded.
NEWS	Print news items. Used by system administrator to keep users informed of current events.
CALENDAR	Automatic reminder service for events of today and tomorrow.
WRITE	Establish direct terminal communication with another user.
WALL	Write to all users.
MESG	Inhibit receipt of messages from WRITE and WALL.
CU	Call up another time-sharing system.
	• Transparent interafce to remote machine.
	• File transmission.
	• Take remote input from local file or put remote output into local file.
	• Remote system need not be UNIX System III.
UUCP	File transfer between CPU's.
	• Automatic queuing until line becomes available and remote machine is up.

- Copy between two remote machines.

- Differences, mail, etc., between two machines.

UUCLEAN   Clean up UUCP spool directory.

UULOG   Maintain a summary log of UUCP and UUX transactions.

UUNAME   List the UUCP names of known systems.

UUSTAT   UUCP status inquiry and job control. Display status of, or cancel, previously specified UUCP commands, or provide general status on UUCP connections to other systems.

UUSUB   Define and monitor a UUCP subnetwork.

UUTO

UUPICK   Public CPU-to-CPU command execution. Gather files from various CPUs, execute a command on a specified CPU, and then send standard output to a file on a specified CPU.

CT   Dial the phone number of a modem attached to a terminal, and spawn a login process to that terminal.

4.1.11 Basic Program Development Tools

Some of these utilities are used as integral parts of the higher level languages described in Section 2.

AR   Maintain archives and libraries. Combines several files into one for housekeeping efficiency.

- Create new archive.

- Update archive by date.

- Replace or delete files.

- Print table of contents.

- Retrieve from archive.

AS.PDP

AS.VAX   Assembler.

- Creates object program consisting of

    code, possibly read-only,

    initialized data or read-write code,

    uninitialized data.

- Relocatable object code is directly executable without further transformation.

- Object code normally includes a symbol table.

- Multiple source files.

- Local labels.

- Conditional assembly.

- "Conditional jump" instructions become branches or branches plus jumps depending on distance.

LIBRARY   The basic run-time library. These routines are used freely by all software.

- Buffered character-by-character I/O.
- Formatted input and output conversion (SCANF and PRINTF) for standard input and output, files, in-memory conversion.
- Storage allocator.
- Time conversions.
- Number conversions.
- Password encryption.
- Quicksort.
- Random number generator.
- Mathematical function library, including trigonometric functions and inverses, exponential, logarithm, square root, Bessel functions.

ADB   Interactive debugger.

- Postmortem dumping.
- Examination of arbitrary files, with no limit on size.
- Interactive breakpoint debugging with the debugger as a separate process.
- Symbolic reference to local and global variables.
- Stack trace for C programs.
- Output formats:

    1-, 2-, or 4-byte integers in octal, decimal, or hex

    single and double floating point

    character and string

    disassembled machine instructions

- Patching.
- Searching for integer, character, or floating patterns.
- Handles separated instruction and data space.

SDB   Symbolic debugger (VAX-11/780 only).

- Dynamic access to program variables and stacks.
- Interactive breakpoint debugging.
- Selective execution of program pieces.
- Symbolic reference to local and global variables.
- Stack tracing.
- Output format control.
- Variable names written in C or F77 format.

OD   Dump any file. Output options include any combination of octal or decimal by words, octal by bytes, ASCII, opcodes, hexadecimal.

● Range of dumping is controllable.

LD   Link edit. Combine relocatable object files. Insert required routines from specified libraries.

      ● Resulting code may be sharable.

      ● Resulting code may have separate instruction and data spaces.

LORDER   Places object file names in proper order for loading, so that files depending on others come after them.

NM   Print the namelist (symbol table) of an object program. Provides control over the style and order of names that are printed.

SIZE   Report the core requirements of one or more object files.

STRIP   Remove the relocation and symbol table information from an object file to save space.

TIME   Run a command and report timing information on it.

PROF   Construct a profile of time spent per routine from statistics gathered by time-sampling the execution of a program. Uses floating point.

      ● Subroutine call frequency and average times for C programs.

MAKE   Controls creation of large programs. Uses a control file specifying source file dependencies to make new version; uses time last changed to deduce minimum amount of work necessary.

      ● Knows about CC, YACC, LEX, etc.

CRASH   Examine system images.

      ● Print location of a process.

      ● Print the user structure of a process.

      ● Generate a kernel stack trace of a process.

      ● Print entries from system tables and buffers.

      ● Print crash dump data and statistics.

XREF   Print a cross reference for C programs.

4.1.12 Programmer's Manual for UNIX System III

MANUAL   Machine-readable version of the Programmer's Manual for UNIX System III.

      ● System overview.

      ● All commands.

      ● All system calls.

      ● All subroutines in C and assembler libraries.

      ● All devices and other special files.

      ● Formats of file system and kinds of files known to system software.

      ● Boot and maintenance procedures.

MAN   Print specified manual section on your terminal.

4.1.13 Error Logging

ERRDEMON	Error-logging daemon. Collects error records from the operating system and places them in a log file.
ERRPT	Process a report of logged errors.
ERRDEAD	Examine a system dump and extract error records.
ERRSTOP	Terminate the error-logging deamon.

4.1.14 Special Device Handler Programs

300	Handle special functions of DASI 300 and 300s terminals.
4014	Paginator for the Textronix 4014 terminal.
450	Handle special functions of the DASI 450 terminal.
HP	Handle special functions of HP 2460 and 2621 - series terminals.
ST	Synchronous terminal control.
VLX	VAX-11/780 LSI console floppy interface.
VPR	Versatec printer spooler.

4.2 Languages

4.2.1 The C Language

CC — Compile and/or link edit programs in the C language. The operating system in UNIX System III, most of the subsystems, and C itself are written in C. For a full description of C, read The C Programming Language, Brian W. Kernighan and Dennis M. Ritchie, Prentice-Hall, 1978.

- General purpose language designed for structured programming.
- Data types include character, integer, float, double, pointers to all types, functions returning above types, arrays of all types, structures and unions of all types.
- Operations intended to give machine-independent control of full machine facility, including to-memory operations and pointer arithmetic.
- Macro preprocessor for parameterized code and inclusion of standard files.
- All procedures recursive, with parameters by value.
- Machine-independent pointer manipulation.
- Object code uses full addressing capability of the PDP-11.
- Runtime library gives access to all system facilities.
- Definable data types.
- Block structure.

LINT — Verifier for C programs. Reports questionable or non- portable usage such as:

Mismatched data declarations and procedure interfaces.

Unused variables, unreachable code, no-effect operations.

Mistyped pointers.

Obsolete syntax.

- Full crossmodule checking of separately compiled programs.

CB   A beautifier for C programs. Does proper indentation and placement of braces.

SCC   C compiler for stand-alone programs. Prepares files for stand-alone execution.

4.2.2 Fortran

F77   A full compiler for ANSI Standard Fortran 77.

- Compatible with C and supporting tools at object level.
- Optional source compatibility with Fortran 66.
- Free format source.
- Optional subscript-range checking, detection of uninitialized variables.
- All widths of arithmetic: 2- and 4-byte integer; 4- and 8-byte real; 8- and 16-byte complex.

RATFOR   Ratfor adds rational control structure a la C to Fortran.

- Compound statements.
- If-else, do, for, while, repeat-until, break, next statements.
- Symbolic constants.
- File insertion.
- Free format source.
- Translation of relationals like >, >=.
- Produces genuine Fortran to carry away.
- May be used with F77.

EFL   Extended Fortran language. Compiles a program written in EFL language into clean Fortran.

- C-like control structures.
- C-like data structures.
- Generic functions.
- Assignment operators (+ = , &= , etc.).
- Logical operators.
- Translation of relationals.
- Free-form imput.
- Defines and includes.

4.2.3 SNOBOL

SNO   A SNOBOL compiler and interpreter. Very similar to SNOBOL 3; its limitations are:

- Function definitions are static.
- Pattern matches are always anchored.
- No built-in functions.

4.2.4 Other Algorithmic Languages

BS   A compiler/interpreter for modest-sized programs. A descendant of Basic and SNOBOL 4 with a little C language thrown in.

- Statements from console execute immediately.
- Statements from a file compile for later execution (by default).
- Line-at-a-time debugging.
- Many builtin functions.

DC   Interactive programmable desk calculator. Has named storage locations as well as conventional stack for holding integers or programs.

- Unlimited precision decimal arithmetic.
- Appropriate treatment of decimal fractions.
- Arbitrary input and output radices, in particular binary, otcal, decimal and hexadecimal.
- Reverse Polish operators:

    + - * /

    remainder, power, square root,

    load, store, duplicate, clear,

    print, enter program text, execute.

BC   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.
- Go-to-less programming.

4.2.5 Macroprocessing

M4   A general purpose macroprocessor.

- Stream-oriented, recognizes macros anywhere in text.
- Syntax fits with functional syntax of most higher-level languages.
- Can evaluate integer arithmetic expressions.

4.2.6 Compiler-compilers

YACC   An LR(1)-based compiler writing system. During execution of resulting parsers, arbitrary C functions may be called to do code generation or semantic actions.

- BNF syntax specifications.

- Precedence relations.

- Accepts formally ambiguous grammars with non-BNF resolution rules.

LEX   Generator of lexical analyzers. Arbitrary C functions may be called upon isolation of each lexical token.

- Full regular expression, plus left and right context dependence.

- Resulting lexical analysers interface cleanly with YACC parsers.

4.3 Text Processing

4.3.1 Document Preparation

ED   Interactive context editor. Random access to all lines of a file.

- Find lines by number or pattern. Patterns may include: specified characters, don't care characters, choices among characters, repetitions of these constructs, beginning of line, end of line.

- Add, delete, change, copy, move, or join lines.

- Permute or split contents of a line.

- Replace one or all instances of a pattern within a line.

- Combine or split files.

- Escape to Shell (command language) during editing.

- Do any of above operations on every pattern-selected line in a given range.

- Optional encryption for extra security.

PTX   Make a permuted (key word in context) index.

SPELL   Look for spelling errors by comparing each word in a document against a word list.

- 25,000-word list includes proper names.

- Handles common prefixes and suffixes.

- Collects words to help tailor local spelling lists.

TYPO   Look for spelling errors by a statistical technique; not limited to English.

CRYPT   Encrypt and decrypt files for security.

HYPHEN   Find hyphenated words.

4.3.2 Document Formatting

TROFF

NROFF   Advanced typesetting. TROFF drives a Graphic Systems phototypesetter; NROFF drives ASCII terminals of all types. TROFF and NROFF accept the same input language.

- Completely definable page format keyed to dynamically planted "interrupts" at specified lines.

- Maintains several separately definable typesetting environments (e.g., one for body text, one for footnotes, and one for unusually elaborate headings).

- Arbitrary number of output pools can be combined at will.

- Macros with substitutable arguments, and macros invocable in mid-line.

- Computation and printing of numerical quantities.

- Conditional execution of macros.

- Tabular layout facility.

- Positions expressible in inches, centimeters, ems, points, machine units or arithmetic combinations thereof.

- Access to character-width computation for unusually difficult layout problems.

- Overstrikes, built-up brackets, horizontal and vertical line drawing.

- Dynamic relative or absolute positioning and size selection, globally or at the character level.

- Can exploit the characteristics of the terminal being used, for approximating special characters, reverse motions, proportional spacing, etc.

The Graphic Systems typesetter has a vocabulary of several 102-character fonts (4 simultaneously) in 15 sizes. TROFF provides terminal output for rough sampling of the product.

NROFF will produce multicolumn output on terminals capable of reverse line feed, or through the postprocessor COL.

High programming skill is required to exploit the formattting capabilities of TROFF and NROFF, although unskilled personnel can easily be trained to enter documents according to canned formats such as those provided by MM, below. TROFF and EQN are essentially identical to NROFF and NEQN so it is usually possible to define interchangeable formats to produce approximate TROFF copy on terminals before actual typesetting. The preprocessor TBL is fully compatible with TROFF and NROFF.

MM   Format documents using the Memorandum Macro (MM) package. Designed for use with NROFF and TROFF. Features include:

- Pagination control.

- Automatically numbered paragraph headings.

- Footnotes.

- List generation.

- Paragraphing, display and indenting.

- Page headers and footers.

- Table of contents.

- Static and Floating displays.

- Two-column output.

- Special formatting macros for preparing memoranda and released papers.

MMT	
MVT	Typeset documents, view graphs, and slides. Similar to MM, except typesetting is via TROFF.
EQN	A mathematical typesetting preprocessor for TROFF. Translates easily readable formulas, either in-line or displayed, into detailed typesetting instructions.

- Automatic calculation of size changes for subscripts, sub-subscripts, etc.

- Full vocabulary of Greek letters and special symbols, such as 'gamma', 'GAMMA', 'integral'.

- Automatic calculation of large bracket sizes.

- 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 nonprogrammers and mathematical typists.

NEQN	A version of EQN for NROFF; accepts the same input language. Prepares formulas for display on any terminal that NROFF knows about, for example, those based on Diablo printing mechanism.

- Same facilities as EQN within graphical capability of terminal.

TBL	A preprocessor for NROFF/TROFF that translates simple descriptions of table layouts and contents into detailed typesetting instructions.

- Computes column widths.

- Handles left- and right-justified columns, centered columns and decimal-point alignment.

- Places column titles.

- Table entries can be text, which is adjusted to fit.

- Can box all or parts of table.

TC	Simulate Graphic Systems typesetter on Tektronix 4014 scope. Useful for checking TROFF page layout before typesetting.
GREEK	Fancy printing on Diablo-mechanism terminals like DASI-300 and DASI-450, and on Tektronix 4014.

- Gives half-line forward and reverse motions.

- Approximates Greek letters and other special characters by overstriking.

COL	Canonicalize files with reverse line feeds for one-pass printing.
DEROFF	Remove all TROFF commands from input.
CHECKEQ	Check document for possible errors in EQN usage.

CW	Prepare TROFF input files for typesetting in the constant-width font.
CHECKCW	Check for possible errors in CW usage.
MMCHEK	Check for possible errors in usage of MM macros and EQN delimiters.

4.4 Information Handling

SORT	Sort or merge ASCII files line-by-line. No limit on input size.

- Sort up or down.
- Sort lexicogaphically or on numeric key.
- Multiple keys located by delimiters or by character position.
- May sort upper case together with lower into dictionary order.
- Optionally suppress duplicate data.

TSORT	Topological sort - converts a partial order into a total order.
UNIQ	Collapse successive duplicate lines in a file into one line.

- Publish lines that were originally unique, duplicated, or both.
- May give redundancy count for each line.

TR	Do one-to-one character translation according to an arbitrary code.

- May coalesce selected repeated characters.
- May delete selected characters.

DIFF	Report line changes, additions and deletions necessary to bring two files into agreement.

- May produce an editor script to convert one file into another.
- A variant compares two new versions against one old one.

BDIFF	Big difference program. Process files too big for DIFF.
DIFF3	3-way differential file comparison. Compares three versions of a file and publishes disagreeing ranges of text.
SDIFF	Side-by-side difference program. Produces a side-by-side listing of two files indicating those lines that are different.
DIFFMK	Compare two versions of a file and create a third file containing "change mark" commands for NROFF or TROFF.
COMM	Identify common lines in two sorted files. Output in up to 3 columns shows lines present in first file only, present in both, and/or present in second only.
JOIN	Combine two files by joining records that have identical keys.
GREP	Print all lines in a file that satisfy a pattern as used in the editor ED.

- May print all lines that fail to match.
- May print count of hits.
- May print first hit in each file.

WC	Count the lines, "words" (blank-separated strings) and characters in a file.

SED	Stream-oriented version of ED. Can perform a sequence of editing operations on each line of an input stream of unbounded lenght.

- Lines may be selected by address or range of addresses.
- Control flow and conditional testing.
- Multiple output streams.
- Multi-line capability.

AWK	Pattern scanning and processing language. Searches input for patterns, and performs actions on each line of input that satisfies the pattern.

- Patterns include regular expressions, arithmetic and lexicographic conditions, boolean combinations and ranges of these.
- Data treated as string or numeric as appropriate.
- Can break input into fields; fields are variables.
- Variables and arrays (with non-numeric subscripts).
- Full set of arithmetic operators and control flow.
- Multiple output streams to files and pipes.
- Output can be formatted as desired.
- Multi-line capabilities.

DIRCMP	Examine two directories and generate tabulated information about their contents.
BFS	Big file scanner. Similar to ED, except it is read-only and processes larger files. Useful for identifying sections of a large file where CSPLIT can be used to divide it.

- Maximum file size is 1024-K bytes.
- Scans actual file, not a copy.
- All ED address expressions are supported.
- Regular expression processing.
- Most ED commands operate.
- Many additional commands.

PWCK	Scan the password file and note inconsistencies.
GRPCK	Verify entries in the group file.
CREF	Make a cross-reference listing of an assembler or C program.

4.5 Other Utility Programs

ARCV	Convert archive files from PDP-11 to VAX-11/780 format.
BCOPY	Interactively copy files starting at arbitrary block boundaries.
FSCV	Convert file systems between PDP-11 and VAX-11/780 formats.
HELP	Print information to explain a message from a command or explain the use of a command.
NL	Line numbering filter. Read lines from a file and prepend line numbers.

REGCMP	Compile regular expressions. Output is C source code.
VC	Version control. Copies lines based on tests of keyword values.
XARGS	Construct an argument list and execute a specified command.

4.6 Graphics

The programs in this section are predominantly intended for use with Tektronix 4014 storage scopes.

GRAPHICS	Establish access to the graphical and numerical commands.
GRAPH	Prepares a graph of a set of input numbers. • Input scaled to fit standard plotting area. • Abscissae may be supplied automatically. • Graph may be labeled. • Control over grid style, line style, graph orientation, etc.
SPLINE	Provides a smooth curve through a set of points intended for GRAPH.
TPLOT	A set of filters for printing graphs produced by GRAPH and other programs on various terminals. Filters provided for Textronix 4014, DASI termainals, Versatec printer/plotter.
GDEV	A set of graphical device routines and filters for the Textronix 4010 series terminals and the Hewlett-Packard 7221A Graphics Plotter.
GED	An interactive graphical editor. Used to display, construct, and edit graphical data files on Textronix 4010 series display terminals.
GUTIL	A set of device independent graphical utility commands.
STAT	A set of command level functions (nodes) that can be interconnected to form a stastical network. Data is passed through the network as sequences of numbers (vectors). There are four classes of nodes: • Transformers map input vector elements into output vector elements. • Summarizers calculate statistics of a vector. • Translators convert among formats. • Generators are sources of definable vectors.
TOC	A set of routines to generate graphical table of contents.

4.7 Source Code Control System (SCCS)

SCCS is a collection of commands for controlling changes to files of text (typically, the source code of programs or text of documents).

ADMIN	Create new SCCS files and change parameters of existing ones.
CDC	Change the delta commentary of an SCCS delta.
COMB	Combine SCCS deltas.
DELTA	Make a delta (change) to an SCCS file.
GET	Create an ASCII text file from a specified SCCS file.
PRS	Print an SCCS file.

RMDEL	Remove a delta from an SCCS file.
SACT	Inform the user of any impending deltas to a named SCCS file.
SCCSDIFF	Compare two versions of an SCCS file and generate a list of differences.
UNGET	Undo a previous GET of an SCCS file.
VAL	Determine if a specified file is an SCCS file meeting characteristics specified by the argument list.
WHAT	Identify SCCS files.

4.8 Remote Job Entry (RJE)

RJE provides for submission and retrieval of jobs from a host system (e.g. an IBM System/360 or System/370). To the host RJE appears to be an IBM 360 remote multileaving work station.

SEND	The command-level interface to the RJE system. It allows the user to collect input from various sources in order to create a run stream consisting of card images, and submit this run stream for transmission to a host computer.
GATH	Concatenate files and write them to the standard output.
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.

### 4.9 Virtual Protocol Machine (VPM)

VPM is a synchronous communications subsystem built around the KMC11 microcomputer.

KAS	An assembler/debugger/loader for the KMC11 microprocessor.
KUN	An un-assembler for the KMC11/DMC11 microprocessor. It produces an output listing, acceptable to the assembler KAS, from object code.
VPMC	A compiler for the protocol description language.

- Generates a load module for the virtual protocol machine.

- Language is essentially a subset of C.

- Uses the RATFOR preprocessor.

- Two versions: BISYNC version, HDLC version.

VPMSTART	Writes load module to the KMC11 and starts the VPM interpreter.
VPMSNAP	Opens the trace drive and reads and prints time-stamped event records.
VPMTRACE	Opens the trace driver and reads and prints event records.

4.10 Novelties, Games, and Things That Didn't Fit Anywhere Else

BACKGAMMON	A player of modest accomplishment.
CHESS	Plays good class D chess.
BJ	A blackjack dealer.
CUBIC	An accomplished player of 4x4x4 tic-tac-toe.
MAZE	Constructs random mazes for you to solve.
MOO	A fascinating number-guessing game.

CAL	Print a calendar of specified month and year.
BANNER	Print output in huge letters.
UNITS	Convert amounts between different scales of measurement. Knows hundreds of units. For example, how many km/sec is a parsec/megayear?
TTT	A tic-tac-toe program that learns. It never makes the same mistake twice.
ARITHMETIC	Speed and accuracy test for number facts.
FACTOR	Factor large integers.
QUIZ	Test your knowledge of Shakespeare, Presidents, capitals, etc.
WUMP	Hunt the wumpus, thrilling search in a dangerous cave.
REVERSI	A two person board game, isomorphic to Othello.
HANGMAN	Word-guessing game. Uses the dictionary supplied with SPELL.
TRUE	
FALSE	Provide truth values.
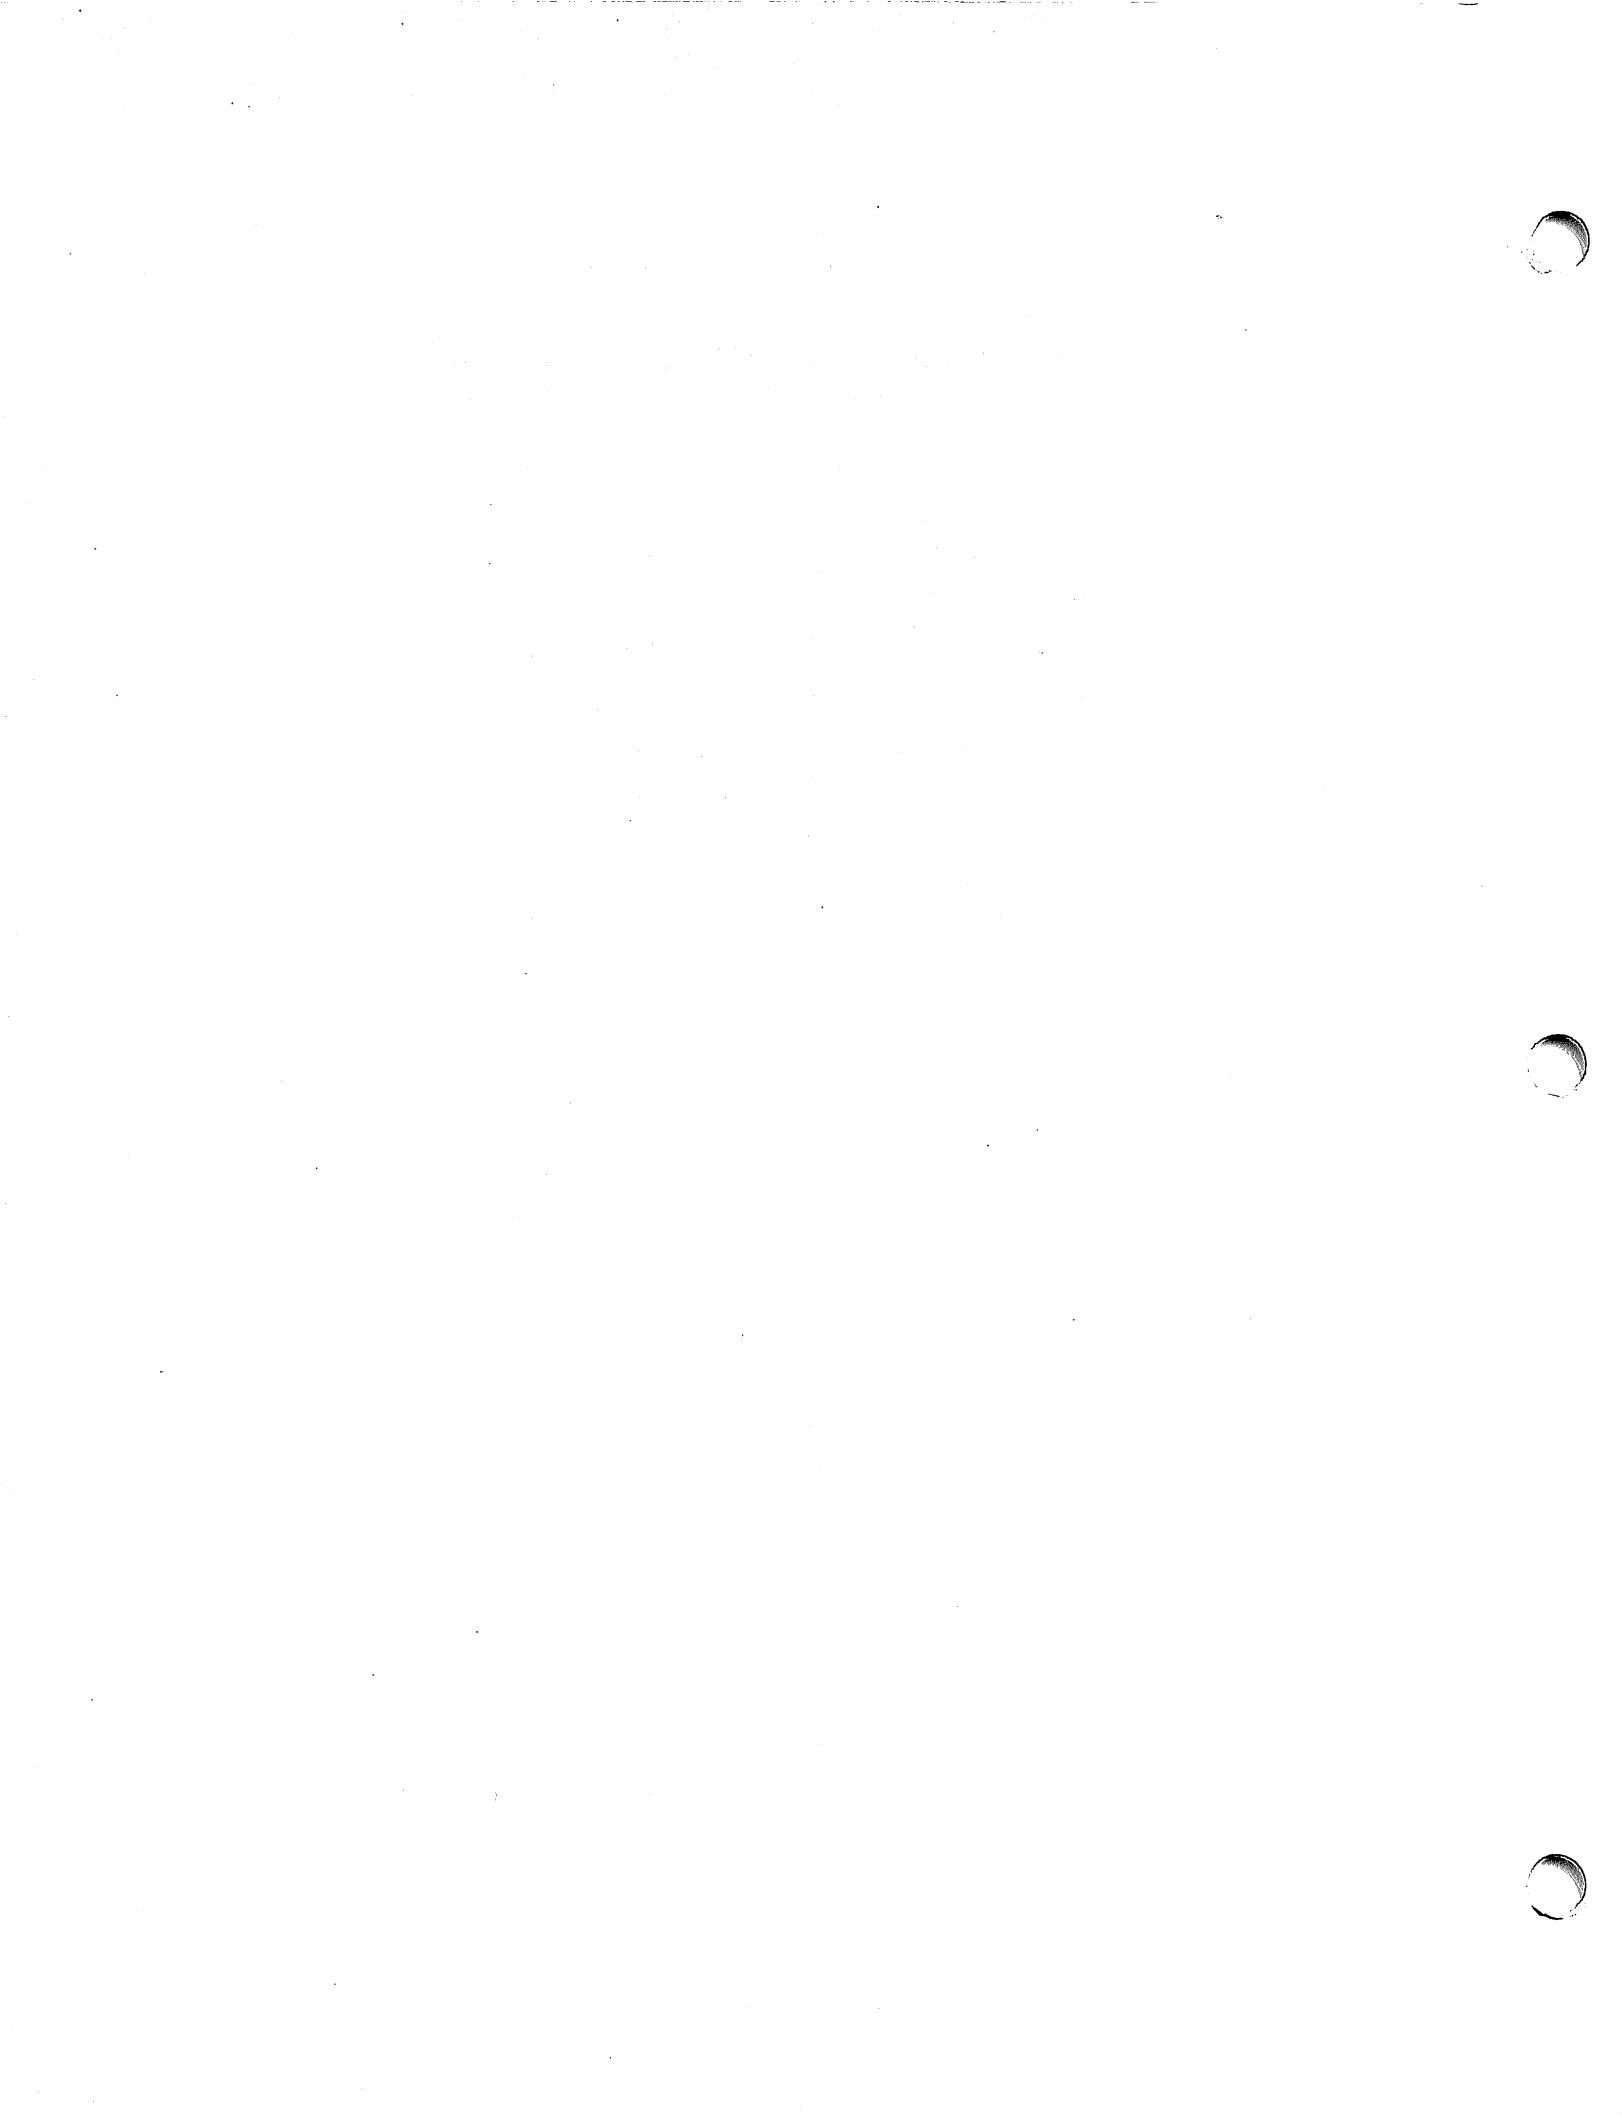
The UNIX Time-Sharing System*

D. M. Ritchie and K. Thompson

ABSTRACT

UNIX† is a general-purpose, multi-user, interactive operating system for the larger Digital Equipment Corporation PDP-11 and the Interdata 8/32 computers. It offers a number of features seldom found even in larger operating systems, including

i   A hierarchical file system incorporating demountable volumes,

ii   Compatible file, device, and inter-process I/O,

iii The ability to initiate asynchronous processes,

iv System command language selectable on a per-user basis,

v   Over 100 subsystems including a dozen languages,

vi High degree of portability.

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

1. INTRODUCTION

There have been four versions of the UNIX time-sharing system. The earliest (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. The third incorporated multiprogramming and ran on the PDP-11/34, /40, /45, /60, and /70 computers; it is the one described in the previously published version of this paper, and is also the most widely used today. This paper describes only the fourth, current system that runs on the PDP-11/70 and the Interdata 8/32 computers. In fact, the differences among the various systems is rather small; most of the revisions made to the originally published version of this paper, aside from those concerned with style, had to do with details of the implementation of the file system.

Since PDP-11 UNIX became operational in February, 1971, over 600 installations have been put into service. Most of them are engaged in applications such as computer science education, the preparation and formatting of documents 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.

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: it can run on hardware costing as little as $40,000, and less than two man-years were spent on the main system software. We hope, however, that users find that the most important

characteristics of the system are its simplicity, elegance, and ease of use.

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

C compiler
Text editor based on QED[1]
Assembler, linking loader, symbolic debugger
Phototypesetting and equation setting programs[2,3]
Dozens of languages including Fortran 77, Basic, Snobol, APL, Algol 68, M6,
TMG, Pascal

There is a host of maintenance, utility, recreation and novelty programs, all written locally. The UNIX user community, which numbers in the thousands, has contributed many more programs and languages. It is worth noting that the system is totally self-supporting. All UNIX software is maintained on the system; likewise, this paper and all other documents in this issue were generated and formatted by the UNIX editor and text formatting programs.

II. HARDWARE AND SOFTWARE ENVIRONMENT

The PDP-11/70 on which the Research UNIX system is installed is a 16-bit word (8-bit byte) computer with 768K bytes of core memory; the system kernel occupies 90K bytes about equally divided between code and data tables. 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 96K bytes of core altogether. There are even larger installations; see the description of the PWB/UNIX systems,[4,5] for example. There are also much smaller, though somewhat restricted, versions of the system.[6]

Our own PDP-11 has two 200-Mb moving-head disks for file system storage and swapping. There are 20 variable-speed communications interfaces attached to 300- and 1200-baud data sets, and an additional 12 communication lines hard-wired to 9600-baud terminals and satellite computers. There are also several 2400- and 4800-baud synchronous communication interfaces used for machine-to-machine file transfer. Finally, there is a variety of miscellaneous devices including nine-track magnetic tape, a line printer, a voice synthesizer, a phototypesetter, a digital switching network, and a chess machine.

The preponderance of UNIX software is written in the abovementioned C language.[7] 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 was about one-third greater than that of the old. Since the new system not only became much easier to understand and to modify but also included many functional improvements, including multiprogramming and the ability to share reentrant code among several user programs, we consider this increase in size quite acceptable.

III. THE FILE SYSTEM

The most important role of the system 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. A file of text consists simply of a string of characters, with lines demarcated by the newline 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 that 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**. Other system directories contain 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 one of these directories 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 that 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. The UNIX system 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 directory, which is its parent. The reason for this is to simplify the writing of programs that 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 supported I/O device 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 as it may to an ordinary file. Thus, for example, to write on a magnetic tape one may write on the file **/dev/mt**. Special files exist for each communication line, each disk, each tape drive, and for physical main memory. Of course, the active disks and the memory 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 with two arguments: the name of an existing ordinary file, and the name of a special file whose associated storage volume (e.g., a 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 a small partition of one of our disk drives, while the other drive, which contains the user's files, is mounted by the system initialization sequence. 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 that would otherwise be required to assure removal of the links whenever the removable volume is dismounted.

3.5 Protection

Although the access control scheme 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 ten protection bits. Nine of these specify independently read, write, and execute permission for the owner of the file, for other members of his group, and for all remaining users.

If the tenth bit is on, the system will temporarily change the user identification (hereafter, user ID) 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 that calls for it. The set-user-ID feature provides for privileged programs that may use files inaccessible to other users. For example, a program may keep an accounting file that should neither be read nor changed except by the program itself. If the set-user-ID 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 that call privileged system entries. For example, there is a system entry invokable only by the "super-user" (below) that 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 "..".

Because 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 by "Aleph-null."[8]

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.

- 5 -

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 number of bytes written on it; no predetermination of the size of a file is necessary or possible.

To illustrate the essentials of I/O, some of the basic calls are summarized below in an anonymous language that will indicate the required parameters without getting into the underlying complexities. 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)

where **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 that 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.

The file system maintains no locks visible to the user, 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 that another user is reading, cannot prevent confusion when, for example, both users are editing a file with an editor that makes a copy of the file being edited.

There are, however, sufficient internal interlocks to maintain the logical consistency of the file system when two users engage simultaneously in activities such 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 immediately following byte. For each open file there is a pointer, maintained inside the system, that 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, such as 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 terminals never return more than one line of input. When a **read** call returns with **n** equal to zero, the end of the file has been reached. 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 terminal by use of an escape sequence that depends on the device used.

Bytes written affect only those parts of a file 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 made to grow 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 = lseek (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 terminals) seek calls are ignored. The actual offset from the beginning of the file to which the pointer was moved is returned in **location**.

There are several additional system entries having to do with I/O and with the file system that 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.

IV. IMPLEMENTATION OF THE FILE SYSTEM

As mentioned in Section 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 found thereby (the file's *i-node*) contains the description of the file:

i   the user and group-ID of its owner

ii   its protection bits

iii   the physical disk or tape addresses for the file contents

iv   its size

v   time of creation, last use, and last modification

vi   the number of links to the file, that is, the number of times it appears in a directory

vii   a code indicating whether the file is a directory, an ordinary file, or a special file.

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, during a subsequent call to read or write the file, the descriptor 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 that 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 de-allocated.

The space on all disks that contain a file system is divided into a number of 512-byte blocks logically addressed from 0 up to a limit that depends on the device. There is space in the i-node of each file for 13 device addresses. For nonspecial files, the first 10 device addresses point at the first 10 blocks of the file. If the file is larger than 10 blocks, the 11 device address points to an indirect block containing up to 128 addresses of additional blocks in the file. Still larger files use the twelfth device address of the i-node to point to a double-indirect block naming 128 indirect blocks, each pointing to 128 blocks of the file. If required, the thirteenth device address is a triple-indirect block. Thus files may conceptually grow to $[(10+128+128^2+128^3) \cdot 512]$ bytes. Once opened, bytes numbered below 5120 can be read with a single disk access; bytes in the range 5120 to 70,656 require two accesses; bytes in the

range 70,656 to 8,459,264 require three accesses; bytes from there to the largest file (1,082,201,088) require four accesses. In practice, a device cache mechanism (see below) proves effective in eliminating most of the indirect fetches.

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 12 device address words are immaterial, and the first specifies an internal *device name*, which is interpreted as a pair of numbers representing, 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 terminal interfaces.

In this environment, the implementation of the **mount** system call (Section 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 that 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 the i-number of the root directory 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; conversely, after a **write** the user's workspace may be reused. In fact, the system maintains a rather complicated buffering mechanism that 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. The system will search its buffers to see whether the affected disk block currently resides in main 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 that reads or writes files in units of 512 bytes has an advantage over a program that reads or writes a single byte at a time, but the gain is not immense; it comes mainly from the avoidance of system overhead. If a program is used rarely or does no great volume of I/O, it 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 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, because 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, because all directory entries for a file have equal status. Charging the owner of a file is unfair in general, for 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. Many installations avoid the issue by not charging any fees at all.

V. PROCESSES AND IMAGES

An *image* is a computer execution environment. It includes a memory 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 main memory; during the execution of other processes it remains in main memory unless the appearance of an active, higher-priority process forces it to be swapped out to the disk.

The user-memory 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 hardware protection 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 stack pointer fluctuates.

5.1 Processes

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

processid = fork ( )

When **fork** is executed, the process splits into two independently executing processes. The two processes have independent copies of the original memory image, and share all open files. The new processes differ only in that one is considered the parent process: in the parent, the returned **processid** actually identifies the child process and is never 0, while in the child, the returned value is always 0.

Because the values returned by **fork** in the parent and child process are distinguishable, each process may determine whether it is the parent or child.

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 Section 6.2), it is not a completely general mechanism, because 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, arg_1, arg_2, ... , arg_n)

which requests the system to read in and execute the program named by **file**, passing it string arguments arg_1, arg_2, ... , arg_n. All the code and data in the process invoking **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 (status)

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. The parent is notified through the **wait** primitive, and **status** is made available to it. Processes may also terminate as a result of various illegal actions or user-generated signals (Section VII below).

VI. THE SHELL

For most users, communication with the system 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. (The shell is described fully elsewhere,[9] so this section will discuss only the theory of its operation.) In simplest form, a command line consists of the command name followed by arguments to the command, all separated by spaces:

command arg_1 arg_2 .... arg_n

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 memory 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 generally prefixes a string such as **/bin/** to **command** and attempts again to find the file. Directory **/bin** contains commands intended to be generally used. (The sequence of directories to be searched may be changed by user request.)

6.1 Standard I/O

The discussion of I/O in Section III 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 three open files with file descriptors 0, 1, and 2. 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 terminal. Thus programs that wish to write informative information ordinarily use file descriptor 1. Conversely, file 0 starts off open for reading, and programs that wish to read messages typed by the user read this file.

The shell is able to change the standard assignments of these file descriptors from the user's terminal 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 keyboard. 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.

File descriptor 2 is, like file 1, ordinarily associated with the terminal output stream. When an output-diversion request with ">" is specified, file 2 remains attached to the terminal, so that commands may produce diagnostic messages that do not silently end up in the output file.

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" requests double-column output.) 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 that we have found useful perform character transliteration, selection of lines according to a pattern, 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 number 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 these examples, an output file other than the terminal 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 &

writes 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 keyboard, **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 also provides general conditional and looping constructions.

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 newline 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 keyboard 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 that 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 memory image of its parent but also all the files currently open in its parent, including those with file descriptors 0, 1, and 2. 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 that are its own standard input and output, because 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 the branch of the return from **fork** belonging to the parent process; that is, the branch that does a **wait,** then reads another command line.) The one thing that 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 <comfile

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. Because this shell process is the child of another instance of the shell, the **wait** executed in the latter will return, and another command may then 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 the system 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 terminal channel. The various subinstances of **init** open the appropriate terminals for input and output on files 0, 1, and 2, waiting, if necessary, for carrier to be established on dial-up lines. Then a message is typed out requesting that the user log in. When the user types a name or other identification, the appropriate instance of **init** wakes up, receives the log-in line, and reads a password file. If the user's 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 protocol is complete.

Meanwhile, the mainstream path of **init** (the parent of all the subinstances of itself that 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 log-in message. Thus a user may log out simply by typing the end-of-file sequence 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, because 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 for.

Recall that after a user has successfully logged in by supplying a 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 log-in 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 might specify that the editor **ed** is to be used instead of the shell. Thus when users of the editing system log in, they are inside the editor and can begin work immediately; also, they can be prevented from invoking 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 the system 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 these games find themselves limited to the game and unable to investigate the (presumably more interesting) offerings of the UNIX system as a whole.

VII. 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. Unless other arrangements have been made, an illegal action causes the system to terminate the process and to write its 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 that are looping, that 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** file. There is also a **quit** signal used to force an image file to be produced. Thus programs that loop unexpectedly may be halted and the remains inspected without prearrangement.

The hardware-generated faults and the interrupt and quit signals can, by request, be either ignored or caught by a 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.

VIII. PERSPECTIVE

Perhaps paradoxically, the success of the UNIX system 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 other author and several colleagues, 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, the system had proved useful enough to persuade management to invest in the PDP-11/45, and later in the PDP-11/70 and Interdata 8/32 machines, upon which it developed to its present form. Our goals throughout the effort, when

articulated at all, have always been to build a comfortable relationship with the machine and to explore ideas and inventions in operating systems and other software. We have not been faced with the need to satisfy someone else's requirements, and for this freedom we are grateful.

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

First: because 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 noninteractive 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 also 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. Because 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, less than a page long, that 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 to push device-dependent considerations into the operating system itself. The only alternatives seem to be to load, with all programs, routines for dealing with each device, 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 the command interface have proved both convenient and efficient. Because 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 that 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.

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 GENIE time-sharing system.[10] On a number of points we were influenced by Multics, which suggested the particular form of the I/O system calls[11] 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 TENEX.[12]

IX. STATISTICS

The following numbers are presented to suggest the scale of the Research UNIX 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 today:

125	user population
33	maximum simultaneous users
1,630	directories
28,300	files
301,700	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, and other semi-infinite problems. Not counting this background work, we average daily:

13,500	commands
9.6	CPU hours
230	connect hours
62	different users
240	log-ins

X. ACKNOWLEDGMENTS

The contributors to UNIX are, in the traditional but here especially apposite phrase, too numerous to mention. Certainly, collective salutes are due to our colleagues in the Computing Science Research Center. R. H. Canaday contributed much to the basic design of the file system. We are particularly appreciative of the inventiveness, thoughtful criticism, and constant support of R. Morris, M. D. McIlroy, and J. F. Ossanna.

References

1. L. P. Deutsch and B. W. Lampson, "An online editor," *Comm. Assoc. Comp. Mach.* **10**(12) pp. 793-799, 803 (December 1967).

2. B. W. Kernighan and L. L. Cherry, "A System for Typesetting Mathematics," *Comm. Assoc. Comp. Mach.* **18** pp. 151-157 (March 1975).

3. B. W. Kernighan, M. E. Lesk, and J. F. Ossanna, "UNIX Time-Sharing System: Document Preparation," *Bell Sys. Tech. J.* **57**(6) pp. 2115-2135 (1978).

4. T. A. Dolotta and J. R. Mashey, "An Introduction to the Programmer's Workbench," *Proc. 2nd Int. Conf. on Software Engineering*, pp. 164-168 (October 13-15, 1976).

5. T. A. Dolotta, R. C. Haight, and J. R. Mashey, "UNIX Time-Sharing System: The Programmer's Workbench," *Bell Sys. Tech. J.* **57**(6) pp. 2177-2200 (1978).

6.   H. Lycklama, "UNIX Time-Sharing System: UNIX on a Microprocessor," *Bell Sys. Tech. J.* 57(6) pp. 2087-2101 (1978).

7.   B. W. Kernighan and D. M. Ritchie, *The C Programming Language,* Prentice-Hall, Englewood Cliffs, New Jersey (1978).

8.   Aleph-null, "Computer Recreations," *Software Practice and Experience* 1(2) pp. 201-204 (April-June 1971).

9.   S. R. Bourne, "UNIX Time-Sharing System: The UNIX Shell," *Bell Sys. Tech. J.* 57(6) pp. 1971-1990 (1978).

10.   L. P. Deutsch and B. W. Lampson, "SDS 930 time-sharing system preliminary reference manual," Doc. 30.10.10, Project GENIE, Univ. Cal. at Berkeley (April 1965).

11.   R. J. Feiertag and E. I. Organick, "The Multics input-output system," *Proc. Third Symposium on Operating Systems Principles,* pp. 35-41 (October 18-20, 1971).

12.   D. G. Bobrow, J. D. Burchfiel, D. L. Murphy, and R. S. Tomlinson, "TENEX, a Paged Time Sharing System for the PDP-10," *Comm. Assoc. Comp. Mach.* 15(3) pp. 135-143 (March 1972).

UNIX Documentation Road Map

G. A. Snyder
J. R. Mashey

Bell Laboratories
Murray Hill, New Jersey 07974

1. INTRODUCTION

A great deal of documentation exists for UNIX.† New users are often overcome by the volume and distributed nature of the documentation. This "road map" attempts to be a terse, up-to-date outline of important documents and information sources.

☛ *The information in this road map is accurate only for UNIX Edition 3.0; it may not apply to other versions of UNIX.*

1.1 Things to Do

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

1.2 Notation Used in This Road Map

{*N* } → Section N in this road map.
•• → Item required for everyone.
• → Item recommended for most users.

All other items are optional and depend on specific interests. Most of the documents mentioned in this road map can be found in *Documents for UNIX*.

Entries in Section *n* of the *UNIX User's Manual* are referred to by *name*(n).

1.3 Prerequisite Structure of Following Sections

```
            {2}
             |
    -----------------
    |   |   |   |
   {3}  {6} {7} {9}
   / \    \ /
 {4} {5}  {8,10}
```
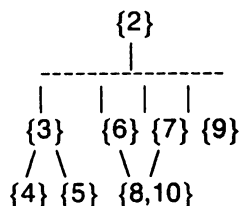
2. BASIC INFORMATION

You won't be able to do much until you have learned most of the material in {2.1}, {2.2}, and {2.3}. You must know how to log into the system, make your terminal work correctly, enter and edit files, and perform basic operations on directories and files.

2.1 UNIX User's Manual ••

- Read *Introduction* and *How to Get Started*.
- Read the *intro* entry in each section.
- Look through Section 1 to become familiar with command names.
- Note the existence of the *Table of Contents* and of the *Permuted Index*.

Section 1 will be especially needed for reference use.

---

† UNIX is a Trademark of Bell Laboratories.

2.2 UNIX for Beginners (Second Edition) ••

2.3 A Tutorial Introduction to the UNIX Text Editor ••

2.4 Advanced Editing on UNIX •

2.5 The Bell System Technical Journal, Vol. 57, No. 6, Part 2 •

Contains several articles on UNIX. In particular, the first paper gives a good overview of UNIX.

2.6 Things to Do

- Do all the exercises found in {2.2} and {2.3}, and maybe {2.4}.
- If you want some sequence of commands to be executed each time you log in, create a file named **.profile** in your login directory.[1] A sample **.profile** can be found in *profile*(5).
- Files in directory **/usr/news** contain recent information on various topics. To print all the news items that have been added since you last looked, type:

        news

2.7 Manual Entries to Be Studied

The following commands are described in Section 1 of the *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).
cd(1)	change working (current) directory.
chmod(1)	change the mode of a file.
cp(1)	copy (*cp*), move (*mv*) or link (*ln*) files.
ed(1)	edit a file.
ls(1)	list a directory; file names beginning with . are not listed unless the -a flag is used.
mkdir(1)	make a (new) directory.
pr(1)	print files (paginated listings).
pwd(1)	print working directory.
rm(1)	remove (delete) file(s); *rmdir* removes the named directories, which must be empty.

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

login(1)	sign on.
mail(1)	send mail to other users or inspect mail from them.
man(1)	print entries of *UNIX User's Manual*.
mesg(1)	permit or deny messages to your terminal.
news(1)	print news items: **news -n** prints a list of recent items.
passwd(1)	change your login password.
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.
term(7)	a list of commonly-used terminals.
who(1)	print list of currently logged-in users.
write(1)	communicate with another (logged-in) user.

Several useful status commands also exist:

date(1)	print time and date.

---

1. The directory you are in right after logging into the system.

du(1)   summarize disk usage.
ps(1)   report active process status.
sum(1)   sum and count blocks in a file.

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 Document Preparation •

This overview of UNIX text processing methods is one of the articles in the BSTJ. (See {2.5} above).

3.5 Manual Entries to Be Studied

mm(1)   print a document using the memorandum macros.
nroff(1)   format or typeset (*troff*) text files; read this to become familiar with options.
spell(1)   identify possible spelling 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:

300(1)   handle special functions of DASI 300, DASI 300s terminals.
4014(1)   paginate output of Tektronix 4014 terminal.
450(1)   handle special functions of DASI 450 (Diablo 1620) terminal.
col(1)   process text for terminals lacking physical reverse vertical motion, such as the Texas Instrument 700 series, Model 43 *TELETYPE®*, etc.
greek(1)   select terminal filter.
hp(1)   handle special functions of Hewlett-Packard 2621, 2640, and 2645 series terminals.
tc(1)   simulate phototypesetter.

4. SPECIALIZED TEXT PROCESSING

The tools listed here 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 also *tbl*(1)).

4.2 Typesetting Mathematics—User's Guide (Second Edition) •

Read this if you need to produce mathematical equations. It describes the use of the equation-setting command *eqn*(1).

4.3 A TROFF Tutorial

An introduction to formatting text with the phototypesetter.

4.4 Manual Entries to Be Studied

cw(1)	use a special constant-width "example" font.
diffmk(1)	mark changes between versions of a file, using output of *diff*(1) to produce "revision bars" in the right margin.
eqn(1)	preprocessor for mathematical equations.
eqnchar(7)	special character definitions for *eqn*(1).
gcat(1)	send phototypesetter output to a central typesetter (site-dependent).
mmt(1)	typeset documents, view graphs, and slides.
tbl(1)	preprocessor for tabular data.

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.

5.1 NROFF/TROFF User's Manual ••

Look at this in detail and try modifying the examples. Read *A TROFF Tutorial* (see {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. Don't re-invent what you can borrow from an existing package (such as PWB/MM-see {3.1} and {3.2}).

5.3 Manual Entries to Be Studied

All entries mentioned in {3.5} and {4.4}.

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 UNIX Shell Tutorial ••

6.3 An Introduction to the UNIX Shell

6.4 Things to Do

If you want to create your own library of commands, for example **/usr/gas/bin**, set the **PATH** parameter in your **.profile** so that your own library is searched when a command is invoked. For example:

    PATH=:$HOME/bin:/bin:/usr/bin

The **HOME** parameter is described in *sh*(1).

6.5 Manual Entries to Be Studied

Read *sh*(1) first; the following entries give further details on commands that are most frequently used within command language programs:

echo(1)	echo arguments (typically to terminal).
env(1)	set environment for command execution.
expr(1)	evaluate an algebraic expression; includes some string operations.
line(1)	read a line from the standard input.
nohup(1)	run a command immune to communications line hang-up.

sh(1)	shell (command interpreter and programming language).
test(1)	evaluate a logical expression.

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 {5}, command language programming {6}, and the commands listed below in {7.4}.

7.1 Things to Do

This road map notes only the most frequently-used commands. It is wise to scan Section 1 of the *UNIX User's Manual* periodically—you will often discover new uses for commands.

7.2 SED—A Non-Interactive Text Editor

7.3 Awk—A Pattern Scanning and Processing Language

7.4 Manual Entries to Be Studied

The starred items below are especially useful for dealing with "fielded data", i.e., data where each line is a sequence of delimiter-separated fields. The following are used to search or edit files in a single pass:

awk(1)*	perform actions on lines matching specified patterns.
grep(1)	search a file for a pattern; more powerful and specialized versions include *egrep* and *fgrep*.
sed(1)*	stream editor.
tr(1)	transliterate (substitute or delete specified characters).

The following compare files in different ways:

cmp(1)	compare files (byte by byte).
comm(1)	print lines common to two files, or lines that appear in only one of the two files.
diff(1)	differential file comparator (minimal editing for conversion).

The following combine files and/or split them apart:

ar(1)	archiver and library maintainer.
cpio(1)	general file copying and archiving.
cut(1)*	cut out selected fields of each line of a file.
paste(1)*	merge lines from several files.
split(1)	split file into chunks of specified size.

The following interrogate files and print information about them:

file(1)	determine file type (best guess).
od(1)	octal dump (and other kinds also).
wc(1)	word (and line and character) count.

Miscellaneous commands:

find(1)	search directory structure for specified kinds of files.
sort(1)*	sort or merge files.
tee(1)	copy single input to several output files.
uniq(1)*	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 The C Programming Language

By B. W. Kernighan and D. M. Ritchie; published by Prentice Hall (1978). Comprehensive text, includes a tutorial and a reference manual. Read the tutorial; try the examples. Check for

updates to the reference manual {8.2} from time to time.

8.2 The C Programming Language-Reference Manual ••

8.3 UNIX Programming •

8.4 YACC—Yet Another Compiler Compiler

8.5 LEX—A Lexical Analyzer Generator

8.6 Lint, a C Program Checker

8.7 Make—A Program for Maintaining Computer Programs

8.8 Things to Do

Read {8.1} and do some of the exercises. A good way to become familiar with 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 **/usr/src**. In particular, the directory **cmd** contains the source for most of the commands. Also, investigate directory **/usr/include**.

8.9 Manual Entries to Be Studied

adb(1)	debug C programs on the PDP-11 series of machines.
cc(1)	compile C programs.
ld(1)	link edit object files; you must know about some of its flags.
lex(1)	generate lexical analyzers.
lint(1)	verify C programs.
make(1)	automate program (re)generation procedures.
nm(1)	print name (i.e., symbol) list.
prof(1)	display profile data; used for program optimization.
ps(1)	report active process status.
sdb(1)	debug C programs symbolically on the VAX 11/780.
strip(1)	remove symbols and relocation bits from executable files.
time(1)	time a command.
yacc(1)	parser generator.

9. NUMERICAL COMPUTATION

9.1 DC—An Interactive Desk Calculator

9.2 BC—An Arbitrary Precision Desk-Calculator Language

9.3 Awk—A Pattern Scanning and Processing Language

9.4 A Portable Fortran 77 Compiler

9.5 Ratfor—A Preprocessor for a Rational Fortran

9.6 Manual Entries to Be Studied

awk(1)	perform actions on lines matching specified patterns.
bc(1)	an interactive language, acts as front end for *dc*(1).
bs(1)	a compiler/interpreter for modest-sized programs.
dc(1)	a desk calculator.
f77(1)	a Fortran compiler.
ratfor(1)	a rational Fortran dialect.

10. INTER-SYSTEM COMMUNICATION

10.1 A Dial-Up Network of UNIX Systems

10.2 Manual Entries to Be Studied

The following commands (most of which are site-dependent) are useful in communicating with other UNIX systems:

cu(1)	call another system.
dpr(1)	print files off-line at a specified destination.
fget(1)	retrieve files from the MH HONEYWELL 6000.
fsend(1)	send files to the MH HONEYWELL 6000.
gcat(1)	send phototypesetter output to the MH HONEYWELL 6000.
uucp(1)	copy files from one UNIX system to another.
uux(1)	execute command(s) on another UNIX system.

May 1979

UNIX For Beginners — Second Edition

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

This paper is meant to help new users get started on the UNIX† operating system. It includes:

- basics needed for day-to-day use of the system — typing commands, correcting typing mistakes, logging in and out, mail, inter-terminal communication, the file system, printing files, redirecting I/O, pipes, and the shell.

- document preparation — a brief discussion of the major formatting programs and macro packages, hints on preparing documents, and capsule descriptions of some supporting software.

- UNIX programming — using the editor, programming the shell, programming in C, other languages and tools.

- An annotated UNIX bibliography.

September 30, 1978

---

† UNIX is a Trademark of Bell Laboratories.

UNIX For Beginners — Second Edition

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

INTRODUCTION

From the user's point of view, the UNIX operating system 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 help new users get used to the main ideas of the UNIX system and start making effective use of it quickly.

You should have a couple of other documents with you for easy reference as you read this one. The most important is *The UNIX Programmer's Manual*; it's often easier to tell you to read about something in the manual than to repeat its contents here. The other useful document is *A Tutorial Introduction to the UNIX Text Editor*, which will tell you how to use the editor to get text — programs, data, documents — into the computer.

A word of warning: the UNIX system has become quite popular, and there are several major variants in widespread use. Of course details also change with time. So although the basic structure of UNIX and how to use it is common to all versions, there will certainly be a few things which are different on your system from what is described here. We have tried to minimize the problem, but be aware of it. In cases of doubt, this paper describes Version 7 UNIX.

This paper has five sections:

1. Getting Started: How to log in, how to type, what to do about mistakes in typing, how to log out. Some of this is dependent on which system 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 the system effectively: generally useful commands; the file system.

3. Document Preparation: Preparing manuscripts is one of the most common uses for UNIX systems. This section contains advice, but not extensive instructions on any of the formatting tools.

4. Writing Programs: UNIX is an excellent system for developing programs. This section talks about some of the tools, but again is not a tutorial in any of the programming languages provided by the system.

5. A UNIX Reading List. An annotated bibliography of documents that new users should be aware of.

1. GETTING STARTED

Logging In

You must have a UNIX login name, which you can get from whoever administers your system. You also need to know the phone number, unless your system uses permanently connected terminals. The UNIX system is capable of dealing with a wide variety of terminals: Terminet 300's; Execuport, TI and similar portables; video (CRT) terminals like the HP2640, etc.; high-priced graphics terminals like the Tektronix 4014; plotting terminals like those from GSI and DASI; and even the venerable Teletype in its various forms. But note: UNIX is strongly oriented towards devices with *lower case*. If your terminal produces only upper case (e.g., model 33 Teletype, some video and portable terminals), life will be so difficult that you should look for another terminal.

Be sure to set the switches appropriately on your device. Switches that might need to be adjusted include the speed, upper/lower case mode, full duplex, even parity, and any others that local wisdom advises. Establish a connection using whatever magic is needed for your terminal; this may involve dialing a telephone call or merely flipping a switch. In either case, UNIX should type "login:" at you. If it types garbage, you may be at the wrong speed; check the switches. If that fails, push the "break" or

"interrupt" key a few times, slowly. 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; the system will not do anything until you type a RETURN. If a password is required, you will be asked for it, and (if possible) printing will be turned off while you type it. Don't forget RETURN.

The culmination of your login efforts is a "prompt character," a single character that indicates that the system is ready to accept commands from you. The prompt character is usually a dollar sign $ or a percent sign %. (You may also get a message of the day just before the prompt character, or a notification that you have mail.)

Typing Commands

Once you've seen the prompt character, you can type commands, which are requests that the system do something. Try typing

> **date**

followed by RETURN. You should get back something like

> **Mon Jan 16 14:17:10 EST 1978**

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. RETURN won't be mentioned again, but don't forget it — it has to be there at the end of each line.

Another command you might try is **who**, which tells you everyone who is currently logged in:

> **who**

gives something like

mb	tty01	Jan 16	09:11
ski	tty05	Jan 16	09:33
gam	tty11	Jan 16	13:07

The time is when the user logged in; "ttyxx" is the system's idea of what terminal the user is on.

If you make a mistake typing the command name, and refer to a non-existent command, you will be told. For example, if you type

> **whom**

you will be told

> **whom: not found**

Of course, if you inadvertently type the name of some other command, it will run, with more or less mysterious results.

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 or a return to the left margin. You can often fix this by logging out and logging back in. Or you can read the description of the command **stty** in section 1 of the manual. To get intelligent treatment of tab characters (which are much used in UNIX) if your terminal doesn't have tabs, type the command

> **stty −tabs**

and the system will convert each tab into the right number of blanks for you. If your terminal 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 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 the line over.

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. So to enter a sharp or at-sign in something, type \# or \@. The system will always echo a newline at you after your at-sign, even if preceded by a backslash. Don't worry — the at-sign has been recorded.

To erase a backslash, you have to type two sharps or two at-signs, as in \##. The backslash is used extensively in UNIX to indicate that the following character is in some way special.

Read-ahead

UNIX has full read-ahead, 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 and interpreted in the correct order. So you can type several 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). The "interrupt" or "break" key found on most terminals can also be used. In a few programs, like the text editor, DEL stops whatever the program is doing but leaves you in that program. Hanging up the phone will stop most programs.

Logging Out

The easiest way to log out is to hang up the phone. You can also type

login

and let someone else use the terminal you were on. It is usually not sufficient just to turn off the terminal. Most UNIX systems do not use a 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 communicate with other users of the system. To read your mail, type the command

mail

Your mail will be printed, one message at a time, most recent message first. After each message, **mail** waits for you to say what to do with it. The two basic responses are **d**, which deletes the message, and RETURN, which does not (so it will still be there the next time you read your mailbox). Other responses are described in the manual. (Earlier versions of **mail** do not process one message at a time, but are otherwise similar.)

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, often called "EOF" for end-of-file, is used throughout the system to mark the end of input from a terminal, so you might as well get used to it.

For practice, send mail to yourself. (This isn't as strange as it might sound — mail to one-self is a handy reminder mechanism.)

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, out of the blue will come a message like

Message from joe tty07...

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 editor tutorial.)

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 (**oo**), 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 **EOF** 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. This is also useful for getting the most up-to-date information on a command. To print a manual section, type "man command-name". Thus to read up on the **who** command, type

man who

and, of course,

man man

tells all about the **man** command.

Computer Aided Instruction

Your UNIX system may have available a program called **learn**, which provides computer aided instruction on the file system and basic commands, the editor, docum nt preparation, and even C programming. Try typing the command

learn

If **learn** exists on your system, it will tell you what to do from there.

II. DAY-TO-DAY USE

Creating Files — The Editor

If you have to type a paper or a letter or a program, how do you 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 called **junk** with some text in it, do the following:

ed junk   (invokes the text editor)
a   (command to "ed", to add text)
now type in
whatever text you want ...
            (signals the end of adding text)

The "." that signals the end of adding text must be at the beginning of a line by itself. Don't forget it, for until it is typed, no other **ed** commands will be recognized — everything you type will be treated as text to be added.

At this point you can do various editing operations on the text you typed in, such as

correcting spelling mistakes, rearranging paragraphs and the like. Finally, you must write the information you have typed into a file with the editor command **w**:

w

ed will respond with the number of characters it wrote into the file **junk**.

Until the **w** command, nothing is stored permanently, so if you hang up and go home the information is lost.† But after **w** the information is there permanently; you can re-access it any time by typing

ed junk

Type a **q** command to quit the editor. (If you try to quit without writing, **ed** will print a **?** to remind you. A second **q** gets you out regardless.)

Now create a second file called **temp** in the same manner. You should now have two files, **junk** and **temp**.

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 you type

ls

the response will be

junk
temp

which are indeed the two files just created. The names are sorted into alphabetical order automatically, but other variations are possible. For example, the command

ls −t

causes the files to be listed 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 Jul 22 2:56 junk
−rw−rw−rw− 1 bwk   78 Jul 22 2:57 temp

The date and time are of the last change to the file. The 41 and 78 are the number of characters (which should agree with the numbers you got from **ed**). **bwk** is the owner of the file, that is, the person who created it. The **−rw−rw−rw−** tells who has permission to read and write the file, in this case everyone.

---
† This is not strictly true — if you hang up while editing, the data you were working on is saved in a file called **ed.hup**, which you can continue with at your next session.

Options can be combined: **ls —lt** gives the same thing as **ls —l**, 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).

The use of optional arguments that begin with a minus sign, like **—t** and **—lt**, is a common convention for UNIX programs. In general, if a program accepts such optional arguments, they precede any filename arguments. It is also vital that you separate the various arguments with spaces: **ls—l** is not the same as **ls —l**.

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
    1,$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 (several thousand 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 prints on the terminal the contents of all the files named in a list. Thus

    cat junk

prints one file, and

    cat junk temp

prints two. The 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 named 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 print **junk** neatly, then skip to the top of a new page and print **temp** neatly.

pr can 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. **pr** has other capabilities as well; see **pr**(1).

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 **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 what equipment is attached to 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 it 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 both of the files named.

You will get a warning message if one of the named files wasn't there, but otherwise **rm**, like most UNIX commands, does its work silently. There is no prompting or chatter, and error messages are occasionally curt. This terseness is sometimes disconcerting to newcomers, but experienced users find it desirable.

What's in a Filename

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 character 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 have already seen, for example, that in the ls command, ls —t means to list in time order. So if you had a file whose name was —t, you would have a tough time listing it by name. Besides the minus sign, there are other characters which have special meaning. To avoid pitfalls, you would do well to use only letters, numbers and the period until you're familiar with the situation.

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 really 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 names of the files in 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 and can occur several times. Thus

> **rm *junk* *temp***

removes all files that contain **junk** or **temp** as any part of their name. As a special case, * by itself matches every filename, so

> **pr ***

prints all your files (alphabetical order), and

> **rm ***

removes *all files*. (You had better be *very* 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. Then you can say

> **pr chap[12349]***

The [...] means to match any of the characters inside the brackets. A range of consecutive letters or digits can be abbreviated, so you can also do this with

> **pr chap[1—49]***

Letters can also be used within brackets: [a—z] matches any character in the range a through z.

The ? pattern matches any single character, so

> **ls ?**

lists all files which have single-character names, and

> **ls —l chap?.1**

lists information about the first file of each chapter (**chap1.1, chap2.1**, etc.).

Of these niceties, * is certainly 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 single quotes, as in

> **ls '?'**

We'll see some more examples of this shortly.

What's in a Filename, Continued

When you first made that file called **junk**, how did the system know that there wasn't another **junk** somewhere else, especially since the person in the next office is also reading this tutorial? The answer is that generally each user has a private *directory*, which contains only the files that belong to him. When you log in, you are "in" your directory. Unless you take special action, when you create a new file, it is made in the directory that you are currently in; this is most often your own directory, and thus the file is unrelated to any other file of the same name that might exist in someone else's directory.

The set of all files is organized into a (usually big) tree, with your files located several branches 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 proper set of branches. Conversely, you can start where you are and walk toward the root.

Let's try the latter first. The basic tools is the command **pwd** ("print working directory"), which prints the name of the directory you are currently in.

Although the details will vary according to the system you are on, if you give the command **pwd**, it will print something like

/usr/your-name

This says that you are currently in the directory **your-name**, which is in turn in the directory **/usr**, which is in turn in the root directory called by convention just **/**. (Even if it's not called **/usr** on your system, you will get something analogous. Make the corresponding changes and read on.)

If you now type

ls /usr/your-name

you should get exactly the same list of file names as you get from a plain **ls**: with no arguments, **ls** lists the contents of the current directory; given the name of a directory, it lists the contents of that directory.

Next, try

ls /usr

This should print a long series of names, among which is your own login name **your-name**. On many systems, **usr** is a directory that contains the directories of all the normal users of the system, like you.

The next step is to try

ls /

You should get a response something like this (although again the details may be different):

bin
dev
etc
lib
tmp
usr

This is a collection of the basic directories of files that the system knows about; we are at the root of the tree.

Now try

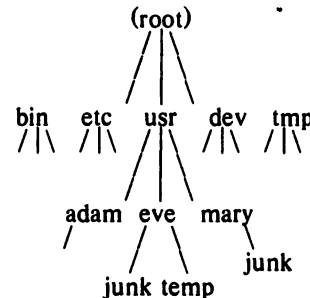**cat /usr/your-name/junk**

(if **junk** is still around in your directory). 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 from the root through the tree of directories to get to a particular file. It is a universal rule in the UNIX system that anywhere you can use an ordinary filename, you can use a pathname.

Here is a picture which may make this clearer:

```
                    (root)
                    /|\
                   / | \
                  /  |  \
            bin  etc  usr  dev  tmp
            /|\  /|\  /|\  /|\  /|\
                      /|\
                     / | \
                    /  |  \
              adam  eve   mary
              /     / \      \
             /     /   \      junk
          junk  temp
```

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 has read-write-execute permissions for the owner, a group, and everyone else, which can be set 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 the prompt character, the system simply looks for a file of that name. It normally 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 a couple of 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:

cd /usr/your-friend

(On some systems, **cd** is spelled **chdir**.) 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. Of course, if you forget what directory you're in, type

pwd

to find out.

It is usually convenient to arrange your own 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

cd book

then start typing chapters. The book is now found in (presumably)

/usr/your-name/book

To remove the directory **book**, type

rm book/*
rmdir book

The first command removes all files from the directory; the second removes the empty directory.

You can go up one level in the tree of files by saying

cd ..

".." 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 systems that the terminal can be replaced by a file for either or both of input and output. As one example,

ls

makes a list of files on your terminal. But if you say

ls >filelist

a list of your files will be placed in the file **filelist** (which will be created if it doesn't already exist, or overwritten if it does). The symbol > means "put the output on the following file, rather than on the terminal." Nothing is produced on the terminal. As another example, you could combine several files into one by capturing the output of **cat** in a file:

cat f1 f2 f3 >temp

The symbol >> operates very much like > does, except that it means "add to the end of." That is,

cat f1 f2 f3 >>temp

means to concatenate **f1**, **f2** and **f3** to the end of whatever is already in **temp**, instead of overwriting the existing contents. As with >, if **temp** doesn't exist, it will be created for you.

In a similar way, 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

As another example, you can use **ed** to prepare a letter in file **let**, then send it to several people with

mail adam eve mary joe <let

Pipes

One of the novel contributions of the UNIX system 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 pipeline.

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** to be neatly formatted.

There are many other examples of pipes. For example,

ls | pr −3

prints a list of your files in three columns. The program **wc** counts the number of lines, words and characters in its input, and as we saw earlier, **who** prints a list of currently-logged on people, one per line. Thus

who | wc

tells how many people are logged on. And of course

ls | wc

counts your files.

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. **pr** is one example:

pr −3 a b c

prints files **a**, **b** and **c** in order in three columns. But in

cat a b c | pr −3

pr prints the information coming down the pipeline, still in three columns.

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, and <, >, and | into changes of input and output streams.

The shell has other capabilities too. For example, you can run 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 prompt character.

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," that is, don't wait for it to complete. 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 say

ed file <script >script.out &

which saves the output lines in a file called **script.out**.

When you initiate a command with **&**, the system 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

If you forget the process number, the command **ps** will tell you about everything you have running. (If you are desperate, **kill 0** will kill all your processes.) And if you're curious about other people, **ps a** will tell you about *all* programs that are currently running.

You can say

(command-1; command-2; command-3) &

to start three 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 simi-

lar 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 **startup,** and then run it with

sh startup

This says to run the shell with the file **startup** as input. The effect is as if you had typed the contents of **startup** on the terminal.

If this is to be a regular thing, you can eliminate the need to type **sh:** simply type, once only, the command

chmod +x startup

and thereafter you need only say

startup

to run the sequence of commands. The **chmod**(1) command marks the file executable; the shell recognizes this and runs it as a sequence of commands.

If you want **startup** to run automatically every time you log in, create a file in your login directory called **.profile,** and place in it the line **startup.** When the shell first gains control when you log in, it looks for the **.profile** file and does whatever commands it finds in it. We'll get back to the shell in the section on programming.

III. DOCUMENT PREPARATION

UNIX systems are used extensively for document preparation. There are two major formatting programs, that is, programs that produce a text with justified right margins, automatic page numbering and titling, automatic hyphenation, and the like. **nroff** is designed to produce output on terminals and line-printers. **troff** (pronounced "tee-roff") instead drives a phototypesetter, which produces very high quality output on photographic paper. This paper was formatted with **troff.**

Formatting Packages

The basic idea of **nroff** and **troff** 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.

Because **nroff** and **troff** are relatively hard to learn to use effectively, several "packages" of canned formatting requests are available to let you specify paragraphs, running titles, footnotes, multi-column output, and so on, with little effort and without having to learn **nroff** and **troff.** These packages take a modest effort to learn, but the rewards for using them are so great that it is time well spent.

In this section, we will provide a hasty look at the "manuscript" package known as −**ms.** Formatting requests typically consist of a period and two upper-case letters, such as .TL, which is used to introduce a title, or .PP to begin a new paragraph.

A document is typed so it looks something like this:

```
.TL
title of document
.AU
author name
.SH
section heading
.PP
paragraph ...
.PP
another paragraph ...
.SH
another section heading
.PP
etc.
```

The lines that begin with a period are the formatting requests. For example, .PP calls for starting a new paragraph. The precise meaning of .PP depends on what output device is being used (typesetter or terminal, for instance), and on what publication the document will appear in. For example, −**ms** normally assumes that a paragraph is preceded by a space (one line in **nroff,** ½ line in **troff**), and the first word is indented. These rules can be changed if you like, but they are changed by changing the interpretation of .PP, not by re-typing the document.

To actually produce a document in standard format using −**ms,** use the command

troff −ms files ...

for the typesetter, and

nroff −ms files ...

for a terminal. The −**ms** argument tells **troff** and **nroff** to use the manuscript package of formatting requests.

There are several similar packages; check with a local expert to determine which ones are in common use on your machine.

Supporting Tools

In addition to the basic formatters, there is a host of supporting programs that help with document preparation. The list in the next few paragraphs is far from complete, so browse through the manual and check with people around you for other possibilities.

eqn and **neqn** let you integrate mathematics into the text of a document, in an easy-to-learn language that closely resembles the way you would speak it aloud. For example, the **eqn** input

sum from i=0 to n x sub i ~=~ pi over 2

produces the output

$$\sum_{i=0}^{n} x_i = \frac{\pi}{2}$$

The program **tbl** provides an analogous service for preparing tabular material; it does all the computations necessary to align complicated columns with elements of varying widths.

refer prepares bibliographic citations from a data base, in whatever style is defined by the formatting package. It looks after all the details of numbering references in sequence, filling in page and volume numbers, getting the author's initials and the journal name right, and so on.

spell and **typo** detect possible spelling mistakes in a document. **spell** works by comparing the words in your document to a dictionary, printing those that are not in the dictionary. It knows enough about English spelling to detect plurals and the like, so it does a very good job. **typo** looks for words which are "unusual", and prints those. Spelling mistakes tend to be more unusual, and thus show up early when the most unusual words are printed first.

grep looks through a set of files for lines that contain a particular text pattern (rather like the editor's context search does, but on a bunch of files). For example,

grep 'ing$' chap*

will find all lines that end with the letters **ing** in the files **chap***. (It is almost always a good practice to put single quotes around the pattern you're searching for, in case it contains characters like ***** or **$** that have a special meaning to the shell.) **grep** is often useful for finding out in which of a set of files the misspelled words detected by **spell** are actually located.

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).

wc counts the words, lines and characters 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 <input >output

sort sorts files in a variety of ways; **cref** makes cross-references; **ptx** makes a permuted index (keyword-in-context listing). **sed** provides many of the editing facilities of **ed**, but can apply them to arbitrarily long inputs. **awk** provides the ability to do both pattern matching and numeric computations, and to conveniently process fields within lines. These programs are for more advanced users, and they are not limited to document preparation. Put them on your list of things to learn about.

Most of these programs are either independently documented (like **eqn** and **tbl**), or are sufficiently simple that the description in the *UNIX Programmer's Manual* is adequate explanation.

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 that 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.

Keep the individual files of a document down to modest size, perhaps ten to fifteen thousand characters. Larger files edit more slowly, and of course if you make a dumb mistake it's better to have clobbered a small file than a big one. Split into files at natural boundaries in the document, for the same reasons that you start each sentence on a new line.

The second aspect of making change easy is to not commit yourself to formatting details too early. One of the advantages of formatting packages like **—ms** is that they permit you to delay decisions to the last possible moment. Indeed, until a document is printed, it is not even decided whether it will be typeset or put on a line printer.

As a rule of thumb, for all but the most trivial jobs, you should type a document in terms of a set of requests like .PP, and then define them appropriately, either by using one of the canned packages (the better way) or by defining your own **nroff** and **troff** commands. 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 commands and request definitions.

IV. PROGRAMMING

There will be no attempt made to teach any of the programming languages available but a few words of advice are in order. One of the reasons why the UNIX system is a productive programming environment is that there is already a rich set of tools available, and facilities like pipes, I/O redirection, and the capabilities of the shell often make it possible to do a job by pasting together programs that already exist instead of writing from scratch.

The Shell

The pipe mechanism lets you fabricate quite complicated operations out of spare parts that already exist. For example, the first draft of the **spell** program was (roughly)

```
cat ...        collect the files
| tr ...       put each word on a new line
| tr ...       delete punctuation, etc.
| sort         into dictionary order
| uniq         discard duplicates
| comm         print words in text
               but not in dictionary
```

More pieces have been added subsequently, but this goes a long way for such a small effort.

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, such as a book, you could laboriously type

```
ed
e chap1.1
1p
$p
e chap1.2
1p
$p
etc.
```

But you can do the job much more easily. One way is to 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. Alternately (and more easily), you can use the fact that the shell will perform loops, repeating a set of commands over and over again for a set of arguments:

```
for i in chap*
do
        ed $i <script
done
```

This sets the shell variable **i** to each file name in turn, then does the command. You can type this command at the terminal, or put it in a file for later execution.

Programming the Shell

An option often overlooked by newcomers is that the shell is itself a programming language, with variables, control flow (**if-else, while, for, case**), subroutines, and interrupt handling. Since there are many building-block programs, you can sometimes avoid writing a new program merely by piecing together some of the building blocks with shell command files.

We will not go into any details here; examples and rules can be found in *An Introduction to the UNIX Shell*, by S. R. Bourne.

Programming in C

If you are undertaking anything substantial, C is the only reasonable choice of programming language: everything in the UNIX system is tuned to it. The system itself is written in C, as are most of the programs that run on it. It is also a easy language to use once you get started. C is introduced and fully described in *The C Programming Language* by B. W. Kernighan and D. M. Ritchie (Prentice-Hall, 1978). Several sections of the manual describe the system interfaces, that is, how you do I/O and similar functions. Read *UNIX Programming* for more complicated things.

Most input and output in C is best handled with the standard I/O library, which provides a set of I/O functions that exist in compatible form on most machines that have C compilers. In general, it's wisest to confine the system interactions in a program to the facilities provided by this library.

C programs that don't depend too much on special features of UNIX (such as pipes) can be moved to other computers that have C compilers. The list of such machines grows daily; in addition to the original PDP-11, it currently

includes at least Honeywell 6000, IBM 370, Interdata 8/32, Data General Nova and Eclipse, HP 2100, Harris /7, VAX 11/780, SEL 86, and Zilog Z80. Calls to the standard I/O library will work on all of these machines.

There are a number of supporting programs that go with C. **lint** checks C programs for potential portability problems, and detects errors such as mismatched argument types and uninitialized variables.

For larger programs (anything whose source is on more than one file) **make** allows you to specify the dependencies among the source files and the processing steps needed to make a new version; it then checks the times that the pieces were last changed and does the minimal amount of recompiling to create a consistent updated version.

The debugger **adb** is useful for digging through the dead bodies of C programs, but is rather hard to learn to use effectively. The most effective debugging tool is still careful thought, coupled with judiciously placed print statements.

The C compiler provides a limited instrumentation service, so you can find out where programs spend their time and what parts are worth optimizing. 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 they are not super accurate or reproducible.

Other Languages

If you *have* to use Fortran, there are two possibilities. 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 **adb, prof**, etc., are all virtually useless with Fortran programs. There may also be a Fortran 77 compiler on your system. If so, this is a viable alternative to Ratfor, and has the non-trivial advantage that it is compatible with C and related programs. (The Ratfor processor and C tools can be used with Fortran 77 too.)

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. The **lex** lexical analyzer generator does the same job for the simpler languages that can be expressed

as regular expressions. It can be used by itself, or as a front end to recognize inputs for a **yacc**-based program. Both **yacc** and **lex** require some sophistication to use, but the initial effort of learning them can be repaid many times over in programs that are easy to change later on.

Most UNIX systems also make available other languages, such as Algol 68, APL, Basic, Lisp, Pascal, and Snobol. Whether these are useful depends largely on the local environment: if someone cares about the language and has worked on it, it may be in good shape. If not, the odds are strong that it will be more trouble than it's worth.

V. UNIX READING LIST

General:

K. L. Thompson and D. M. Ritchie, *The UNIX Programmer's Manual*, Bell Laboratories, 1978. 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 need to read section 1.

Documents for Use with the UNIX Time-sharing System. Volume 2 of the Programmer's Manual. This contains more extensive descriptions of major commands, and tutorials and reference manuals. All of the papers listed below are in it, as are descriptions of most of the programs mentioned above.

D. M. Ritchie and K. L. Thompson, "The UNIX Time-sharing System," 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.

The Bell System Technical Journal (BSTJ) Special Issue on UNIX, July/August, 1978, contains many papers describing recent developments, and some retrospective material.

The 2nd International Conference on Software Engineering (October, 1976) contains several papers describing the use of the Programmer's Workbench (PWB) version of UNIX.

Document Preparation:

B. W. Kernighan, "A Tutorial Introduction to the UNIX Text Editor" and "Advanced Editing on UNIX," Bell Laboratories, 1978. Beginners need the introduction; the advanced material will help you get the most out of the editor.

M. E. Lesk, "Typing Documents on UNIX," Bell Laboratories, 1978. Describes the **−ms** macro package, which isolates the novice from the vagaries of **nroff** and **troff**, and takes care of

most formatting situations. If this specific package isn't available on your system, something similar probably is. The most likely alternative is the PWB/UNIX macro package —mm; see your local guru if you use PWB/UNIX.

B. W. Kernighan and L. L. Cherry, "A System for Typesetting Mathematics," Bell Laboratories Computing Science Tech. Rep. 17.

M. E. Lesk, "Tbl — A Program to Format Tables," Bell Laboratories CSTR 49, 1976.

J. F. Ossanna, Jr., "NROFF/TROFF User's Manual," Bell Laboratories CSTR 54, 1976. troff is the basic formatter used by —ms, eqn and tbl. The reference manual is indispensable if you are going to write or maintain these or similar programs. But start with:

B. W. Kernighan, "A TROFF Tutorial," Bell Laboratories, 1976. An attempt to unravel the intricacies of troff.

Programming:

B. W. Kernighan and D. M. Ritchie, *The C Programming Language,* Prentice-Hall, 1978. Contains a tutorial introduction, complete discussions of all language features, and the reference manual.

B. W. Kernighan and D. M. Ritchie, "UNIX Programming," Bell Laboratories, 1978. Describes how to interface with the system from C programs: I/O calls, signals, processes.

S. R. Bourne, "An Introduction to the UNIX Shell," Bell Laboratories, 1978. An introduction and reference manual for the Version 7 shell. Mandatory reading if you intend to make effective use of the programming power of this shell.

S. C. Johnson, "Yacc — Yet Another Compiler-Compiler," Bell Laboratories CSTR 32, 1978.

M. E. Lesk, "Lex — A Lexical Analyzer Generator," Bell Laboratories CSTR 39, 1975.

S. C. Johnson, "Lint, a C Program Checker," Bell Laboratories CSTR 65, 1977.

S. I. Feldman, "MAKE — A Program for Maintaining Computer Programs," Bell Laboratories CSTR 57, 1977.

J. F. Maranzano and S. R. Bourne, "A Tutorial Introduction to ADB," Bell Laboratories CSTR 62, 1977. An introduction to a powerful but complex debugging tool.

S. I. Feldman and P. J. Weinberger, "A Portable Fortran 77 Compiler," Bell Laboratories, 1978. A full Fortran 77 for UNIX systems.
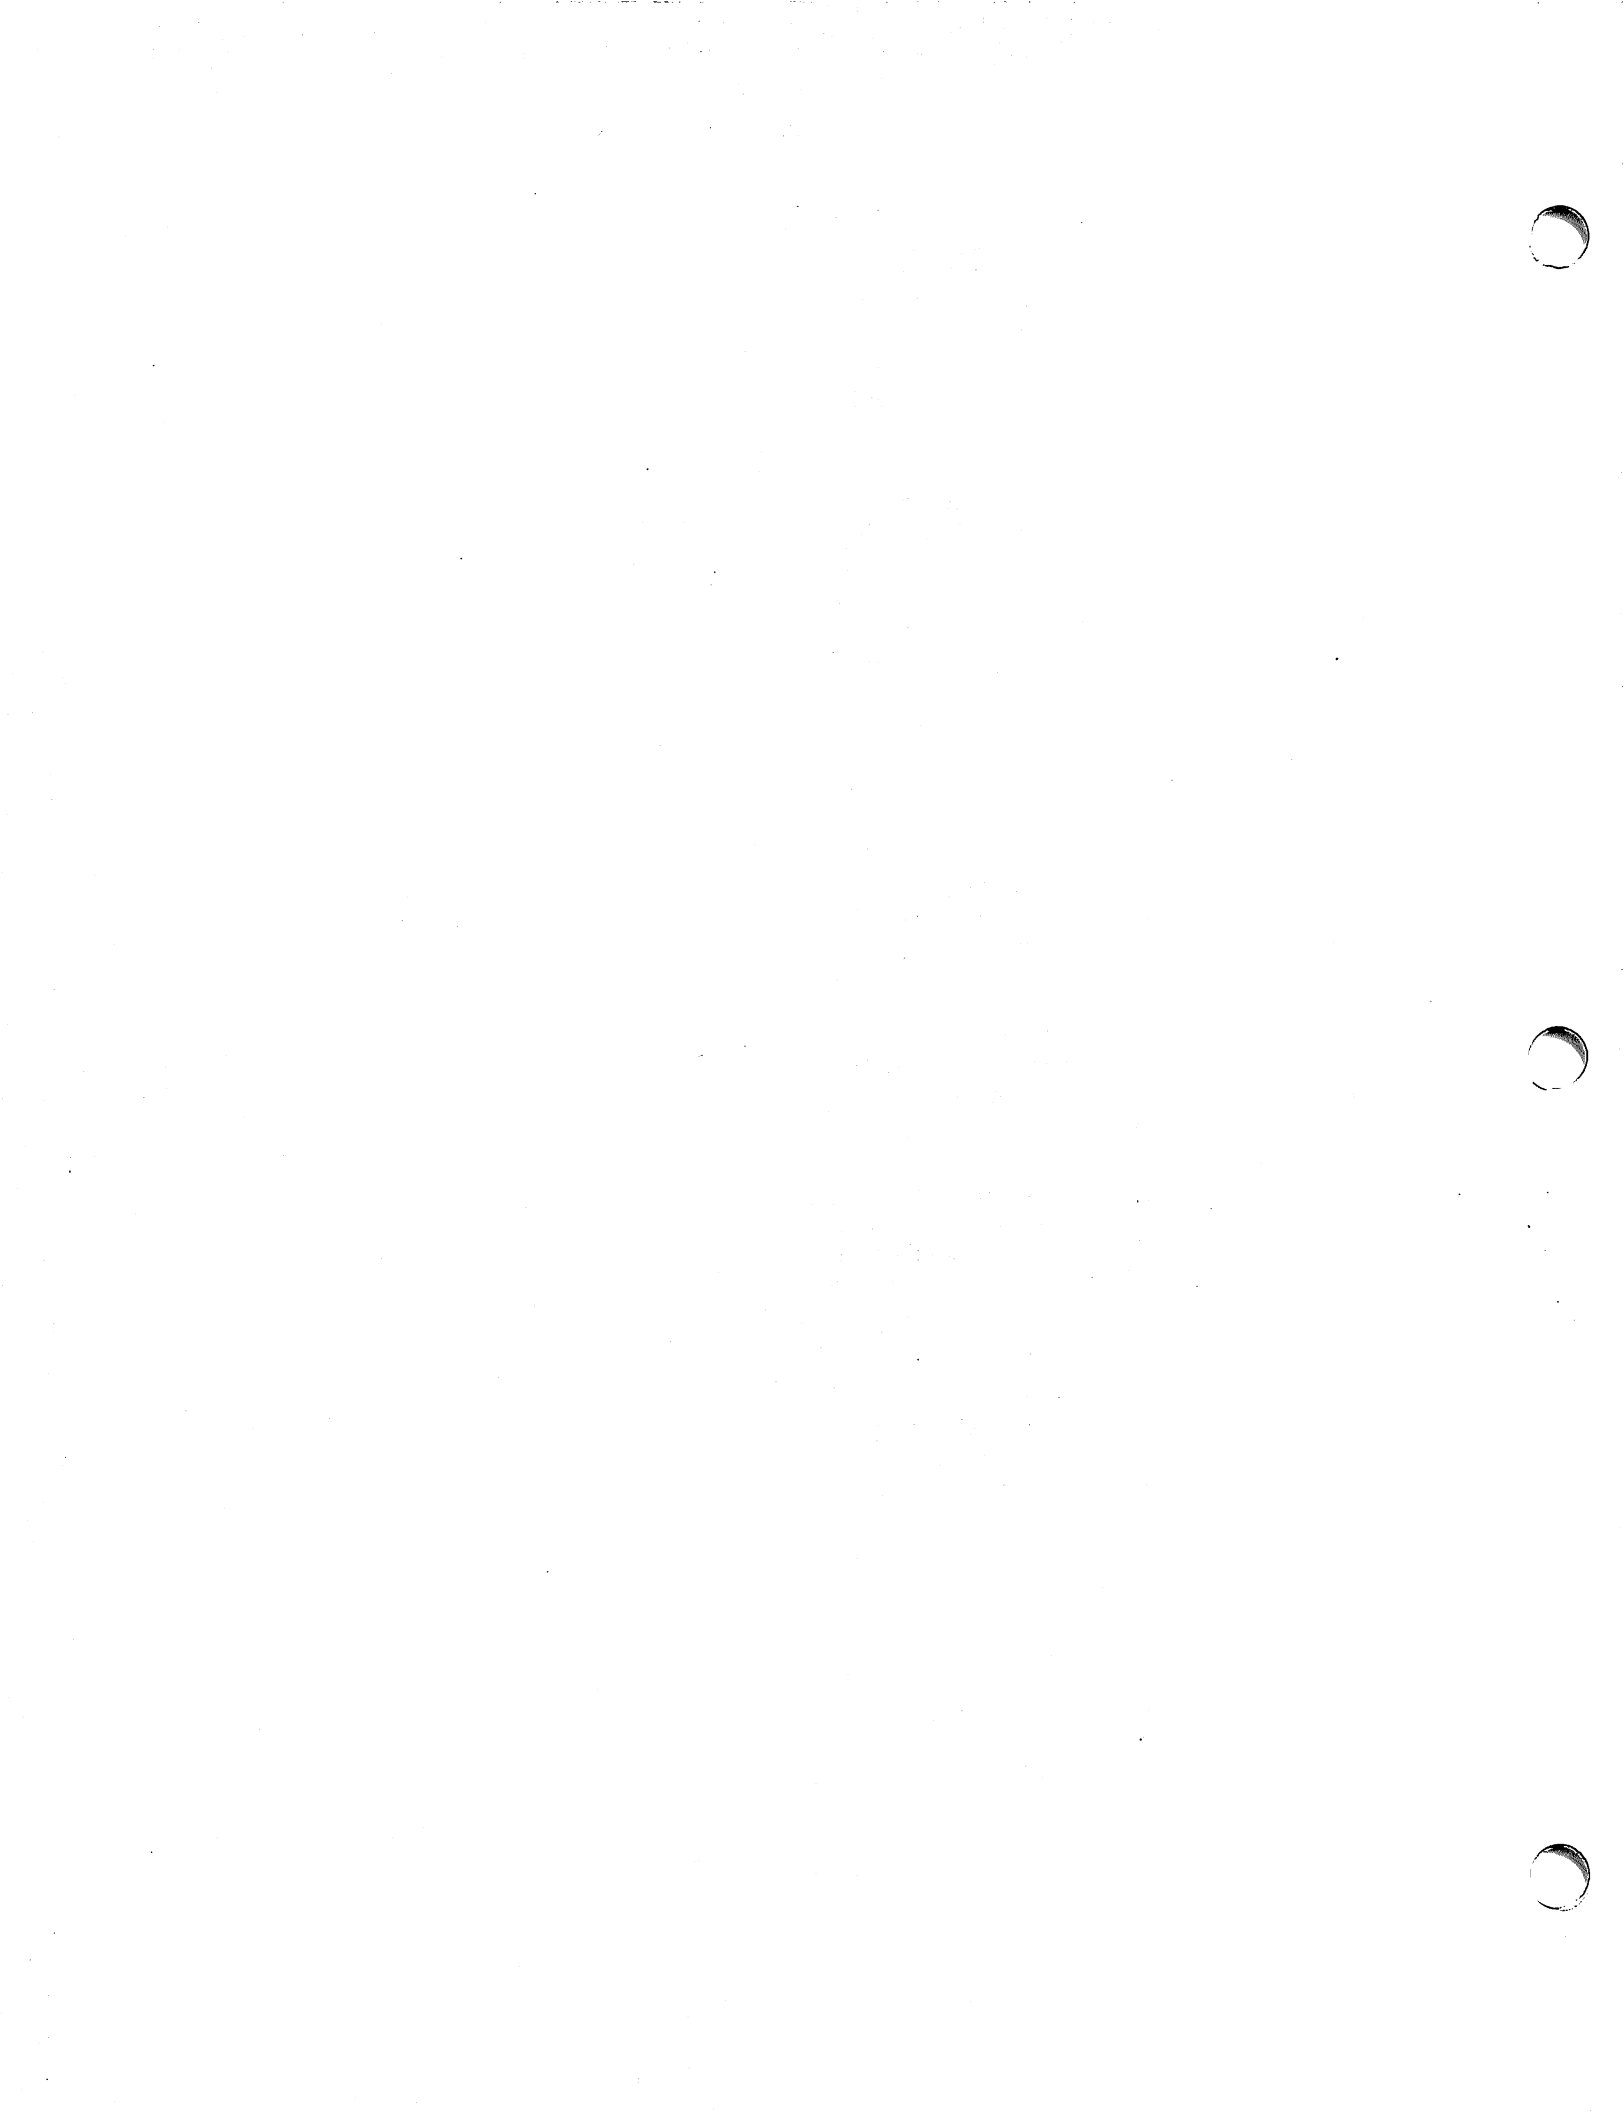
A Tutorial Introduction to the UNIX Text Editor

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

Almost all text input on the UNIX† operating system is done with the text-editor *ed.* This memorandum is a tutorial guide to help beginners get started with text editing.

Although it does not cover everything, it does discuss enough for most users' day-to-day needs. This includes printing, appending, changing, deleting, moving and inserting entire lines of text; reading and writing files; context searching and line addressing; the substitute command; the global commands; and the use of special characters for advanced editing.

September 21, 1978

---

†UNIX is a Trademark of Bell Laboratories.

A Tutorial Introduction to the UNIX Text Editor

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

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 Programmer's 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). When you have mastered the Tutorial, try *Advanced Editing on UNIX*. 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. For more on that, read *UNIX for Beginners*.

You must also know what character to type as the end-of-line on your particular terminal. This character is the RETURN key on most terminals. Throughout, we will refer to this character, whatever it is, as RETURN.

Getting Started

We'll assume that you have logged in to your system and it has just printed the prompt character, usually either a $ or a %. The easiest way to get *ed* is to type

    ed   **(followed by a return)**

You are now ready to go — *ed* is waiting for you to tell it what to do.

Creating Text — the Append command "a"

As your first problem, suppose you want to create some text starting from scratch. Perhaps you 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, just type an a followed by a RETURN, followed by

the lines of text you 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 you 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 you 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 you'll want to save your text for later use. To write out the contents of the buffer onto a file, use the *write* command

        w

followed by the filename you 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 this case, *ed* would respond with

        68

(Remember that blanks and the return 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 you 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 the prompt character ($ or %). 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 the prompt character. 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 you 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

---

† Actually, *ed* will print ? if you try to quit without writing. At that point, write if you want; if not, another q will get you out regardless.

68

which is the number of characters in **junk**. *If anything was already in the buffer, it is deleted first.*

If you use the e command to read a file into the buffer, then you 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 good way to operate is

> ed
> e file
> [editing session]
> w
> q

This way, you can simply say w from time to time, and be secure in the knowledge that if you got the file name right at the beginning, you are writing into the proper file each time.

You can find out at any time what file name *ed* is remembering by typing the *file* command f. In this example, if you typed

> f

ed would reply

> junk

Reading text from a file — the Read command "r"

Sometimes you 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 you 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 **?name**, where **name** is the name of a file; this means that the file doesn't exist, typically because you spelled the file name wrong, or perhaps that you are not allowed to read or write it. 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, use the print command

> p

The way this is done is as follows. Specify the lines where you want printing to begin and where you 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) say

> 1,2p   (starting line = 1, ending line = 2 p)

Ed will respond with

> Now is the time
> for all good men

Suppose you want to print *all* the lines in the buffer. You could use 1,3p as above if you knew there were exactly 3 lines in the buffer. But in general, you don't know how many there are, so what do you 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, you could use

$,$p

but *ed* lets you abbreviate this to

$p

You 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 you abbreviate even further: you can print any single line by typing *just* the line number — no need to type the letter **p**. So if you say

$

ed will print the last line of the buffer.

You can also use $ in combinations like

$−1,$p

which prints the last two lines of the buffer. This helps when you want to see how far you got in typing.

Exercise 3:

As before, create some text using the **a** 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 your buffer still contains the six lines as above, that you have just typed

1,3p

and *ed* has printed the three lines for you. 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 you have done anything with. (You just printed it!) You 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 you did anything to (in this case, line 3, which you 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 you most recently did something to." You 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 example these are lines 3 through 6.

Some commands change the value of dot, while others do not. The **p** command sets dot to the number of the last line printed; the last command will set both . and $ to 6.

Dot is most useful when used in combinations like this one:

.+1   (or equivalently, .+1p)

This means "print the next line" and is a handy way to step slowly through a buffer. You can also say

.−1   (or .−1p)

which means "print the line *before* the current line." This enables you to go backwards if you 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 return will cause printing of the next line — it's equivalent to .+1p. Try it. Try typing a −; you will find that it's equivalent to .−1p.

Deleting lines: the "d" command

Suppose you 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 you can check by using

1,$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

0r *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 you 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*. You 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) say

p

and get

Now is the time

which is what you 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 you 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, you can try again. (Notice that there is a **p** on the same line as the **s** command. 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 you had

Nowxx is the time

you can say

s/xx//p

to get

Now is the time

Notice that // (two adjacent slashes) 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, you can move on to another highly important idea of *ed* — context searching.

Suppose you have the original three line text in the buffer:

Now is the time
for all good men
to come to the aid of their party.

Suppose you want to find the line that contains *their* so you 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 you'd been making changes, deleting and rearranging lines, and so on, you 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 to 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* command

/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.

You 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. They were used 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, you could say

> /Now/ +1s/good/bad/

or

> /good/s/good/bad/

or

> /party/ −1s/good/bad/

The choice is dictated only by convenience. You 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 you don't know how many lines are involved. (Of course, if there were only three lines in the buffer, you'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 a.)

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 you want to put the first three lines of the buffer at the end instead. You could do it by saying:

> 1,3w temp
> $r temp
> 1,3d

(Do you see why?) but you 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 you had

> **First paragraph**
> . . .
> **end of first paragraph.**
> **Second paragraph**
> . . .
> **end of second paragraph.**

you could reverse the two paragraphs like this:

> /Second/,/end of 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/B
> .+2s/ghi/jkl/B
> .−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,* . means "any character," not a period, so

/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,

/[0123456789]/

matches any single digit — any one of the characters inside the braces will cause a match. This can be abbreviated to [0—9].

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 you wanted to put parentheses around it. You could just retype the line, but this is tedious. Or you could say

s/^/(/
s/$/)/

using your knowledge of ^ and $. But the easiest way uses the &:

s/.*/(&)/

This says "match the whole line, and replace it by itself surrounded by parentheses." The & can be used several times in a line; consider using

s/.*/&? &!!/

to produce

Now is the time? Now is the time!!

You don't have to match the whole line, of course: if the buffer contains

the end of the world

you 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 **&** saves you from typing it again.

The **&** is a special character only within the replacement text of a substitute command, and has no special meaning elsewhere. You 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, that is, 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 the specified lines to the new text which follows. The new lines are terminated by a **.,** as with **a.** If no lines are specified, replace line dot. Dot is set to last line changed.

d: 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 new file. Any previous contents of the buffer are thrown away, so issue a **w** beforehand.

f: Print remembered filename. If a name follows **f** the remembered name will be set to it.

g: The command

> **g/---/commands**

will execute the commands on those lines that contain ---, which can be any context search expression.

i: Insert lines before specified line (or dot) until a **.** is typed on a new line. Dot is set to last line inserted.

m: Move lines specified to after the line named after **m.** Dot is set to the last line moved.

p: Print specified lines. If none specified, print line dot. A single line number is equivalent to *line-number* **p.** A single return prints **.+1,** the next line.

q: Quit *ed.* Wipes out all text in buffer if you give it twice in a row without first giving a **w** command.

r: Read a file into buffer (at end unless specified elsewhere.) Dot set to last line read.

s: The command

> **s/string1/string2/**

substitutes the characters **string1** into **string2** in the specified lines. If no lines are specified, make the 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: The command

> **v/---/commands**

executes **commands** on those lines that *do not* contain ---.

w: Write out buffer onto a file. Dot is not changed.

.=: Print value of dot. (**=** by itself prints the value of **$.**)

!: The line

> **!command-line**

causes **command-line** to be executed as a UNIX command.

/-----/: Context search. Search for next line which contains this string of characters. Print it. Dot is set to the line where string was 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 users with any background should find helpful hints on how to get their jobs done more easily.

August 4, 1978

---

† UNIX is a Trademark of Bell Laboratories.

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,

---

†UNIX is a Trademark of Bell Laboratories.

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

    1,$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 > and each backspace as <. 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 \, 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 \07 or \16. 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 particular 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 visible.)

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 con-

tains 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

\x\y

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 the following 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 '◻*' 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', in other words, one or more spaces.

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

[0123456789]

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.

It's a nuisance to have to spell out the digits, so you can abbreviate them as [0−9]; similarly, [a−z] stands for the lower case letters, and [A−Z] for upper case.

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 '^':

[^0−9]

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.

Joining Lines

Lines may also be joined together, but this is done with the **j** command instead of **s**. Given the lines

Now is
□the time

and supposing that dot is set to the first of them,

then the command

.j

joins them together. No blanks are added, which is why we carefully showed a blank at the beginning of the second line.

All by itself, a **j** command joins line dot to line dot+1, but any contiguous set of lines can be joined. Just specify the starting and ending line numbers. For example,

1,$jp

joins all the lines into one big one and prints it. (More on line numbers in Section 3.)

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 sub-

stitute 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 0r 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 got changed. If there were no changes, 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. (Don't look too hard for 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?1

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 lines beginning with '.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'; the system 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' any more), 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* 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 the system 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 the system. 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. (As an aside, if you examine the sequence of commands here, you can see why many UNIX systems use e as a synonym for **ed**.)

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. If you are positioned on the second line,

```
m− −
```

does the interchange.

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 be a ghastly 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 really big, it may be impossible because of limits in **ed**.

The program **grep** was invented to get around these limitations. The search 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 contains '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. 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.
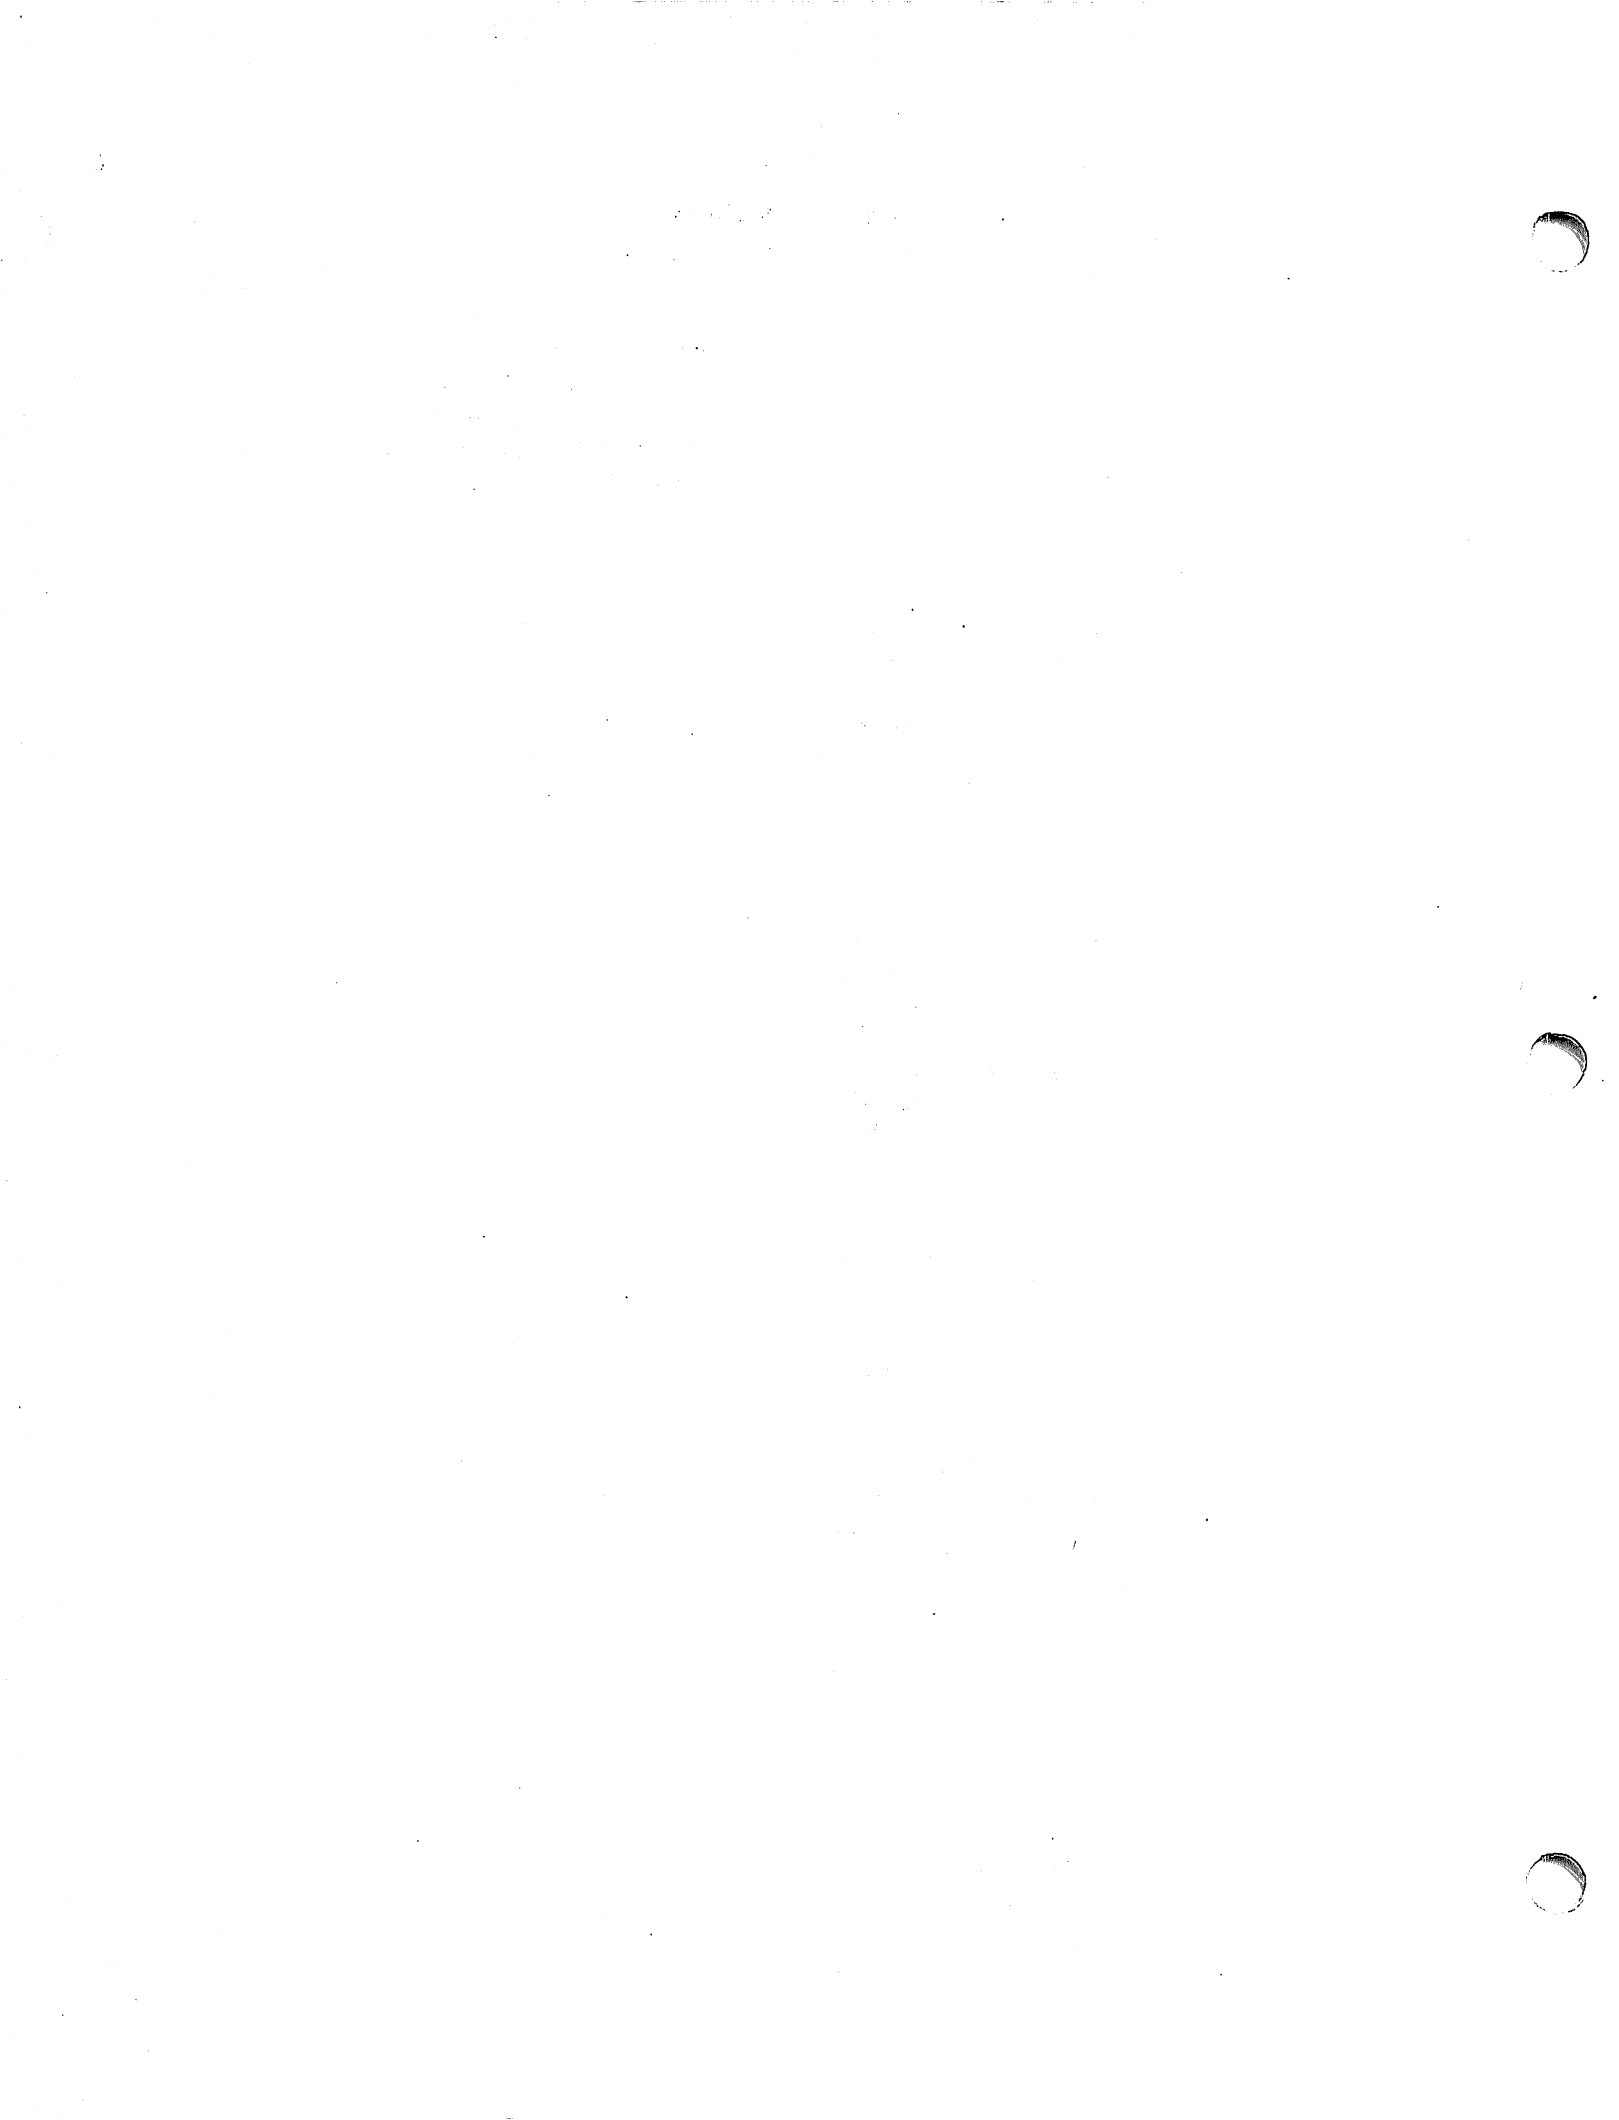
An Introduction to the UNIX Shell

S. R. Bourne

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

The *shell* is a command programming language that provides an interface to the UNIX† operating system. Its features include control-flow primitives, parameter passing, variables and string substitution. Constructs such as *while, if then else, case* and *for* are available. Two-way communication is possible between the *shell* and commands. String-valued parameters, typically file names or flags, may be passed to a command. A return code is set by commands that may be used to determine control-flow, and the standard output from a command may be used as shell input.

The *shell* can modify the environment in which commands run. Input and output can be redirected to files, and processes that communicate through 'pipes' can be invoked. Commands are found by searching directories in the file system in a sequence that can be defined by the user. Commands can be read either from the terminal or from a file, which allows command procedures to be stored for later use.

November 12, 1978

---

†UNIX is a Trademark of Bell Laboratories.

An Introduction to the UNIX Shell

S. R. Bourne

Bell Laboratories
Murray Hill, New Jersey 07974

1.0 Introduction

The shell is both a command language and a programming language that provides an interface to the UNIX operating system. This memorandum describes, with examples, the UNIX shell. The first section covers most of the everyday requirements of terminal users. Some familiarity with UNIX is an advantage when reading this section; see, for example, "UNIX for beginners".[1] Section 2 describes those features of the shell primarily intended for use within shell procedures. These include the control-flow primitives and string-valued variables provided by the shell. A knowledge of a programming language would be a help when reading this section. The last section describes the more advanced features of the shell. References of the form "see *pipe* (2)" are to a section of the UNIX manual.[2]

1.1 Simple commands

Simple commands consist of one or more words separated by blanks. The first word is the name of the command to be executed; any remaining words are passed as arguments to the command. For example,

    who

is a command that prints the names of users logged in. The command

    ls −l

-prints a list of files in the current directory. The argument −*l* tells *ls* to print status information, size and the creation date for each file.

1.2 Background commands

To execute a command the shell normally creates a new *process* and waits for it to finish. A command may be run without waiting for it to finish. For example,

    cc pgm.c &

calls the C compiler to compile the file *pgm.c*. The trailing & is an operator that instructs the shell not to wait for the command to finish. To help keep track of such a process the shell reports its process number following its creation. A list of currently active processes may be obtained using the *ps* command.

1.3 Input output redirection

Most commands produce output on the standard output that is initially connected to the terminal. This output may be sent to a file by writing, for example,

    ls −l >file

The notation >*file* is interpreted by the shell and is not passed as an argument to *ls*. If *file* does not exist then the shell creates it; otherwise the original contents of *file* are replaced with the output from *ls*. Output may be appended to a file using the notation

> ls −l >> file

In this case *file* is also created if it does not already exist.

The standard input of a command may be taken from a file instead of the terminal by writing, for example,

> wc < file

The command *wc* reads its standard input (in this case redirected from *file*) and prints the number of characters, words and lines found. If only the number of lines is required then

> wc −l < file

could be used.

1.4 Pipelines and filters

The standard output of one command may be connected to the standard input of another by writing the 'pipe' operator, indicated by |, as in,

> ls −l | wc

Two commands connected in this way constitute a *pipeline* and the overall effect is the same as

> ls −l > file; wc < file

except that no *file* is used. Instead the two processes are connected by a pipe (see *pipe* (2)) and are run in parallel. Pipes are unidirectional and synchronization is achieved by halting *wc* when there is nothing to read and halting *ls* when the pipe is full.

A *filter* is a command that reads its standard input, transforms it in some way, and prints the result as output. One such filter, *grep*, selects from its input those lines that contain some specified string. For example,

> ls | grep old

prints those lines, if any, of the output from *ls* that contain the string *old*. Another useful filter is *sort*. For example,

> who | sort

will print an alphabetically sorted list of logged in users.

A pipeline may consist of more than two commands, for example,

> ls | grep old | wc −l

prints the number of file names in the current directory containing the string *old*.

1.5 File name generation

Many commands accept arguments which are file names. For example,

> ls −l main.c

prints information relating to the file *main.c*.

The shell provides a mechanism for generating a list of file names that match a pattern. For example,

> ls −l *.c

generates, as arguments to *ls*, all file names in the current directory that end in *.c*. The character * is a pattern that will match any string including the null string. In general *patterns* are specified as follows.

   *   Matches any string of characters including the null string.

   ?   Matches any single character.

  [...]   Matches any one of the characters enclosed. A pair of characters separated by a minus will match any character lexically between the pair.

For example,

      [a—z]*

matches all names in the current directory beginning with one of the letters *a* through *z*.

      /usr/fred/test/?

matches all names in the directory **/usr/fred/test** that consist of a single character. If no file name is found that matches the pattern then the pattern is passed, unchanged, as an argument.

This mechanism is useful both to save typing and to select names according to some pattern. It may also be used to find files. For example,

      echo /usr/fred/*/core

finds and prints the names of all *core* files in sub-directories of **/usr/fred** . (*echo* is a standard UNIX command that prints its arguments, separated by blanks.) This last feature can be expensive, requiring a scan of all sub-directories of **/usr/fred** .

There is one exception to the general rules given for patterns. The character '.' at the start of a file name must be explicitly matched.

      echo *

will therefore echo all file names in the current directory not beginning with '.'.

      echo .*

will echo all those file names that begin with '.'. This avoids inadvertent matching of the names '.' and '..' which mean 'the current directory' and 'the parent directory' respectively. (Notice that *ls* suppresses information for the files '.' and '..'.)

1.6 Quoting

Characters that have a special meaning to the shell, such as < > * ? | & , are called metacharacters. A complete list of metacharacters is given in appendix B. Any character preceded by a \ is *quoted* and loses its special meaning, if any. The \ is elided so that

      echo \?

will echo a single ? , and

      echo \\

will echo a single \. To allow long strings to be continued over more than one line the sequence \newline is ignored.

\ is convenient for quoting single characters. When more than one character needs quoting the above mechanism is clumsy and error prone. A string of characters may be quoted by enclosing the string between single quotes. For example,

      echo xx'****'xx

will echo

      xx****xx

The quoted string may not contain a single quote but may contain newlines, which are preserved. This quoting mechanism is the most simple and is recommended for casual use.

A third quoting mechanism using double quotes is also available that prevents interpretation of some but not all metacharacters. Discussion of the details is deferred to section 3.4.

1.7 Prompting

When the shell is used from a terminal it will issue a prompt before reading a command. By default this prompt is '$'. It may be changed by saying, for example,

    PS1=yesdear

that sets the prompt to be the string *yesdear*. If a newline is typed and further input is needed then the shell will issue the prompt '> '. Sometimes this can be caused by mistyping a quote mark. If it is unexpected then an interrupt (DEL) will return the shell to read another command. This prompt may be changed by saying, for example,

    PS2=more

1.8 The shell and login

Following *login* (1) the shell is called to read and execute commands typed at the terminal. If the user's login directory contains the file **.profile** then it is assumed to contain commands and is read by the shell before reading any commands from the terminal.

1.9 Summary

- **ls**
  Print the names of files in the current directory.

- **ls >file**
  Put the output from *ls* into *file*.

- **ls | wc −1**
  Print the number of files in the current directory.

- **ls | grep old**
  Print those file names containing the string *old*.

- **ls | grep old | wc −1**
  Print the number of files whose name contains the string *old*.

- **cc pgm.c &**
  Run *cc* in the background.

2.0 Shell procedures

The shell may be used to read and execute commands contained in a file. For example,

        sh file [args ...]

calls the shell to read commands from *file*. Such a file is called a *command procedure* or *shell procedure*. Arguments may be supplied with the call and are referred to in *file* using the positional parameters $1, $2, ... For example, if the file *wg* contains

        who | grep $1

then

        sh wg fred

is equivalent to

        who | grep fred

UNIX files have three independent attributes, *read, write* and *execute*. The UNIX command *chmod* (1) may be used to make a file executable. For example,

        chmod +x wg

will ensure that the file *wg* has execute status. Following this, the command

        wg fred

is equivalent to

        sh wg fred

This allows shell procedures and programs to be used interchangeably. In either case a new process is created to run the command.

As well as providing names for the positional parameters, the number of positional parameters in the call is available as $#. The name of the file being executed is available as $0.

A special shell parameter $* is used to substitute for all positional parameters except $0. A typical use of this is to provide some default arguments, as in,

        nroff −T450 −ms $*

which simply prepends some arguments to those already given.

2.1 Control flow - for

A frequent use of shell procedures is to loop through the arguments ($1, $2, ...) executing commands once for each argument. An example of such a procedure is *tel* that searches the file /usr/lib/telnos that contains lines of the form

        ...
        fred mh0123
        bert mh0789
        ...

The text of *tel* is

        for i
        do grep $i /usr/lib/telnos; done

The command

        tel fred

prints those lines in /usr/lib/telnos that contain the string *fred*.

```
tel fred bert
```

prints those lines containing *fred* followed by those for *bert*.

The **for** loop notation is recognized by the shell and has the general form

```
for name in w1 w2 ...
do command-list
done
```

A *command-list* is a sequence of one or more simple commands separated or terminated by a newline or semicolon. Furthermore, reserved words like **do** and **done** are only recognized following a newline or semicolon. *name* is a shell variable that is set to the words *w1 w2* ... in turn each time the *command-list* following **do** is executed. If **in** *w1 w2* ... is omitted then the loop is executed once for each positional parameter; that is, **in** *$\$*$* is assumed.

Another example of the use of the **for** loop is the *create* command whose text is

```
for i do >$i; done
```

The command

```
create alpha beta
```

ensures that two empty files *alpha* and *beta* exist and are empty. The notation >*file* may be used on its own to create or clear the contents of a file. Notice also that a semicolon (or newline) is required before **done**.

2.2 Control flow - case

A multiple way branch is provided for by the **case** notation. For example,

```
case $# in
    1) cat >>$1 ;;
    2) cat >>$2 <$1 ;;
    *) echo 'usage: append [ from ] to' ;;
esac
```

is an *append* command. When called with one argument as

```
append file
```

$\$#$ is the string *1* and the standard input is copied onto the end of *file* using the *cat* command.

```
append file1 file2
```

appends the contents of *file1* onto *file2*. If the number of arguments supplied to *append* is other than 1 or 2 then a message is printed indicating proper usage.

The general form of the **case** command is

```
case word in
    pattern ) command-list ;;
    ...
esac
```

The shell attempts to match *word* with each *pattern*, in the order in which the patterns appear. If a match is found the associated *command-list* is executed and execution of the **case** is complete. Since * is the pattern that matches any string it can be used for the default case.

A word of caution: no check is made to ensure that only one pattern matches the case argument. The first match found defines the set of commands to be executed. In the example below the commands following the second * will never be executed.

```
case $# in
    *) ... ;;
    *) ... ;;
esac
```

Another example of the use of the **case** construction is to distinguish between different forms of an argument. The following example is a fragment of a *cc* command.

```
for i
do case $i in
    -[ocs])      ... ;;
    -*) echo 'unknown flag $i' ;;
    *.c) /lib/c0 $i ... ;;
    *)   echo 'unexpected argument $i' ;;
    esac
done
```

To allow the same commands to be associated with more than one pattern the **case** command provides for alternative patterns separated by a | . For example,

```
case $i in
    -x | -y)    ...
esac
```

is equivalent to

```
case $i in
    -[xy])      ...
esac
```

The usual quoting conventions apply so that

```
case $i in
    \?)         ...
```

will match the character ? .

2.3 Here documents

The shell procedure *tel* in section 2.1 uses the file /usr/lib/telnos to supply the data for *grep*. An alternative is to include this data within the shell procedure as a *here* document, as in,

```
for i
do grep $i <<!
    ...
    fred mh0123
    bert mh0789
    ...
!
done
```

In this example the shell takes the lines between <<! and ! as the standard input for *grep*. The string ! is arbitrary, the document being terminated by a line that consists of the string following <<.

Parameters are substituted in the document before it is made available to *grep* as illustrated by the following procedure called *edg*.

```
ed $3 <<%
g/$1/s//$2/g
w
%
```

The call

```
edg string1 string2 file
```

is then equivalent to the command

```
ed file <<%
g/string1/s//string2/g
w
%
```

and changes all occurrences of *string1* in *file* to *string2*. Substitution can be prevented using \ to quote the special character $ as in

```
ed $3 <<+
1,\$s/$1/$2/g
w
+
```

(This version of *edg* is equivalent to the first except that *ed* will print a ? if there are no occurrences of the string $1.) Substitution within a *here* document may be prevented entirely by quoting the terminating string, for example,

```
grep $i <<\#
...
#
```

The document is presented without modification to *grep*. If parameter substitution is not required in a *here* document this latter form is more efficient.

2.4 Shell variables

The shell provides string-valued variables. Variable names begin with a letter and consist of letters, digits and underscores. Variables may be given values by writing, for example,

```
user=fred box=m000 acct=mh0000
```

which assigns values to the variables **user, box** and **acct.** A variable may be set to the null string by saying, for example,

```
null=
```

The value of a variable is substituted by preceding its name with $; for example,

```
echo $user
```

will echo *fred.*

Variables may be used interactively to provide abbreviations for frequently used strings. For example,

```
b=/usr/fred/bin
mv pgm $b
```

will move the file *pgm* from the current directory to the directory **/usr/fred/bin .** A more general notation is available for parameter (or variable) substitution, as in,

```
echo ${user}
```

which is equivalent to

echo $user

and is used when the parameter name is followed by a letter or digit. For example,

tmp=/tmp/ps
ps a >${tmp}a

will direct the output of *ps* to the file **/tmp/psa**, whereas,

ps a >$tmpa

would cause the value of the variable **tmpa** to be substituted.

Except for $? the following are set initially by the shell. $? is set after executing each command.

$?   The exit status (return code) of the last command executed as a decimal string. Most commands return a zero exit status if they complete successfully, otherwise a non-zero exit status is returned. Testing the value of return codes is dealt with later under **if** and **while** commands.

$#   The number of positional parameters (in decimal). Used, for example, in the *append* command to check the number of parameters.

$$ The process number of this shell (in decimal). Since process numbers are unique among all existing processes, this string is frequently used to generate unique temporary file names. For example,

ps a >/tmp/ps$$
...
rm /tmp/ps$$

$!   The process number of the last process run in the background (in decimal).

$–   The current shell flags, such as −x and −v.

Some variables have a special meaning to the shell and should be avoided for general use.

$MAIL   When used interactively the shell looks at the file specified by this variable before it issues a prompt. If the specified file has been modified since it was last looked at the shell prints the message *you have mail* before prompting for the next command. This variable is typically set in the file **.profile**, in the user's login directory. For example,

MAIL=/usr/mail/fred

$HOME   The default argument for the *cd* command. The current directory is used to resolve file name references that do not begin with a /, and is changed using the *cd* command. For example,

cd /usr/fred/bin

makes the current directory **/usr/fred/bin**.

cat wn

will print on the terminal the file *wn* in this directory. The command *cd* with no argument is equivalent to

cd $HOME

This variable is also typically set in the the user's login profile.

$PATH   A list of directories that contain commands (the *search path*). Each time a command is executed by the shell a list of directories is searched for an executable

file. If $PATH is not set then the current directory, **/bin**, and **/usr/bin** are searched by default. Otherwise $PATH consists of directory names separated by **:**. For example,

PATH = :/usr/fred/bin :/bin :/usr/bin

specifies that the current directory (the null string before the first **:**), **/usr/fred/bin**, **/bin** and **/usr/bin** are to be searched in that order. In this way individual users can have their own 'private' commands that are accessible independently of the current directory. If the command name contains a **/** then this directory search is not used; a single attempt is made to execute the command.

$PS1   The primary shell prompt string, by default, '$ '.

$PS2   The shell prompt when further input is needed, by default, '> '.

$IFS   The set of characters used by *blank interpretation* (see section 3.4).

2.5 The test command

The *test* command, although not part of the shell, is intended for use by shell programs. For example,

test −f file

returns zero exit status if *file* exists and non-zero exit status otherwise. In general *test* evaluates a predicate and returns the result as its exit status. Some of the more frequently used *test* arguments are given here, see *test* (1) for a complete specification.

test s	true if the argument *s* is not the null string
test −f file	true if *file* exists
test −r file	true if *file* is readable
test −w file	true if *file* is writable
test −d file	true if *file* is a directory

2.6 Control flow - while

The actions of the **for** loop and the **case** branch are determined by data available to the shell. A **while** or **until** loop and an **if then else** branch are also provided whose actions are determined by the exit status returned by commands. A **while** loop has the general form

while *command-list₁*
do *command-list₂*
done

The value tested by the **while** command is the exit status of the last simple command following **while**. Each time round the loop *command-list₁* is executed; if a zero exit status is returned then *command-list₂* is executed; otherwise, the loop terminates. For example,

while test $1
do ...
    shift
done

is equivalent to

for i
do ...
done

shift is a shell command that renames the positional parameters $2, $3, ... as $1, $2, ... and loses $1 .

Another kind of use for the **while/until** loop is to wait until some external event occurs and then run some commands. In an **until** loop the termination condition is reversed. For example,

>until test −f file
>do sleep 300; done
>*commands*

will loop until *file* exists. Each time round the loop it waits for 5 minutes before trying again. (Presumably another process will eventually create the file.)

2.7 Control flow - if

Also available is a general conditional branch of the form,

>**if** *command-list*
>**then** *command-list*
>**else** *command-list*
>**fi**

that tests the value returned by the last simple command following **if**.

The **if** command may be used in conjunction with the *test* command to test for the existence of a file as in

>if test −f file
>then *process file*
>else *do something else*
>fi

An example of the use of **if, case** and **for** constructions is given in section 2.10.

A multiple test **if** command of the form

>if ...
>then ...
>else if ...
>   then ...
>   else if ...
>   ...
>   fi
>   fi
>fi

may be written using an extension of the **if** notation as,

>if ...
>then ...
>elif ...
>then ...
>elif ...
>...
>fi

The following example is the *touch* command which changes the 'last modified' time for a list of files. The command may be used in conjunction with *make* (1) to force recompilation of a list of files.

```
flag=
for i
do case $i in
    -c) flag=N ;;
    *)  if test -f $i
        then    ln $i junk$$; rm junk$$
        elif test $flag
        then    echo file \'$i\' does not exist
        else    >$i
        fi
    esac
done
```

The −c flag is used in this command to force subsequent files to be created if they do not already exist. Otherwise, if the file does not exist, an error message is printed. The shell variable *flag* is set to some non-null string if the −c argument is encountered. The commands

```
ln ...; rm ...
```

make a link to the file and then remove it thus causing the last modified date to be updated. The sequence

```
if command1
then    command2
fi
```

may be written

```
command1 && command2
```

Conversely,

```
command1 | | command2
```

executes *command2* only if *command1* fails. In each case the value returned is that of the last simple command executed.

2.8 Command grouping

Commands may be grouped in two ways,

```
{ command-list ; }
```

and

```
( command-list )
```

In the first *command-list* is simply executed. The second form executes *command-list* as a separate process. For example,

```
(cd x; rm junk )
```

executes *rm junk* in the directory x without changing the current directory of the invoking shell. The commands

```
cd x; rm junk
```

have the same effect but leave the invoking shell in the directory x.

2.9 Debugging shell procedures

The shell provides two tracing mechanisms to help when debugging shell procedures. The first is invoked within the procedure as

    set −v

(v for verbose) and causes lines of the procedure to be printed as they are read. It is useful to help isolate syntax errors. It may be invoked without modifying the procedure by saying

    sh −v proc ...

where *proc* is the name of the shell procedure. This flag may be used in conjunction with the −n flag which prevents execution of subsequent commands. (Note that saying *set −n* at a terminal will render the terminal useless until an end-of-file is typed.)

The command

    set −x

will produce an execution trace. Following parameter substitution each command is printed as it is executed. (Try these at the terminal to see what effect they have.) Both flags may be turned off by saying

    set −

and the current setting of the shell flags is available as $− .

2.10 The man command

The following is the *man* command which is used to print sections of the UNIX manual. It is called, for example, as

    man sh
    man −t ed
    man 2 fork

In the first the manual section for *sh* is printed. Since no section is specified, section 1 is used. The second example will typeset (−t option) the manual section for *ed*. The last prints the *fork* manual page from section 2.

```
cd /usr/man

: 'colon is the comment command'
: 'default is nroff ($N), section 1 ($s)'
N=n s=1

for i
do case $i in

    [1—9]*)     s=$i ;;

    —t) N=t ;;

    —n) N=n ;;

    —*) echo unknown flag \'$i\' ;;

    *)  if test —f man$s/$i.$s
        then    ${N}roff man0/${N}aa man$s/$i.$s
        else    : 'look through all manual sections'
                found=no
                for j in 1 2 3 4 5 6 7 8 9
                do if test —f man$j/$i.$j
                    then man $j $i
                        found=yes
                    fi
                done
                case $found in
                    no) echo '$i: manual page not found'
                esac
        fi
    esac
done
```

Figure 1. A version of the man command

3.0 Keyword parameters

Shell variables may be given values by assignment or when a shell procedure is invoked. An argument to a shell procedure of the form *name=value* that precedes the command name causes *value* to be assigned to *name* before execution of the procedure begins. The value of *name* in the invoking shell is not affected. For example,

    user=fred command

will execute *command* with **user** set to *fred*. The −k flag causes arguments of the form *name=value* to be interpreted in this way anywhere in the argument list. Such *names* are sometimes called keyword parameters. If any arguments remain they are available as positional parameters $1, $2, ...

The *set* command may also be used to set positional parameters from within a procedure. For example,

    set − *

will set $1 to the first file name in the current directory, $2 to the next, and so on. Note that the first argument, −, ensures correct treatment when the first file name begins with a −.

3.1 Parameter transmission

When a shell procedure is invoked both positional and keyword parameters may be supplied with the call. Keyword parameters are also made available implicitly to a shell procedure by specifying in advance that such parameters are to be exported. For example,

    export user box

marks the variables **user** and **box** for export. When a shell procedure is invoked copies are made of all exportable variables for use within the invoked procedure. Modification of such variables within the procedure does not affect the values in the invoking shell. It is generally true of a shell procedure that it may not modify the state of its caller without explicit request on the part of the caller. (Shared file descriptors are an exception to this rule.)

Names whose value is intended to remain constant may be declared *readonly*. The form of this command is the same as that of the *export* command,

    readonly name ...

Subsequent attempts to set readonly variables are illegal.

3.2 Parameter substitution

If a shell parameter is not set then the null string is substituted for it. For example, if the variable **d** is not set

    echo $d

or

    echo ${d}

will echo nothing. A default string may be given as in

    echo ${d−.}

which will echo the value of the variable **d** if it is set and '.' otherwise. The default string is evaluated using the usual quoting conventions so that

    echo ${d−'*'}

will echo * if the variable **d** is not set. Similarly

echo ${d−$1}

will echo the value of **d** if it is set and the value (if any) of **$1** otherwise. A variable may be assigned a default value using the notation

echo ${d=.}

which substitutes the same string as

echo ${d−.}

and if **d** were not previously set then it will be set to the string '.'. (The notation ${...=...} is not available for positional parameters.)

If there is no sensible default then the notation

echo ${d?message}

will echo the value of the variable **d** if it has one, otherwise *message* is printed by the shell and execution of the shell procedure is abandoned. If *message* is absent then a standard message is printed. A shell procedure that requires some parameters to be set might start as follows.

: ${user?} ${acct?} ${bin?}
...

Colon (:) is a command that is built in to the shell and does nothing once its arguments have been evaluated. If any of the variables **user, acct** or **bin** are not set then the shell will abàndon execution of the procedure.

3.3 Command substitution

The standard output from a command can be substituted in a similar way to parameters. The command *pwd* prints on its standard output the name of the current directory. For example, if the current directory is **/usr/fred/bin** then the command

d=`pwd`

is equivalent to

d=/usr/fred/bin

The entire string between grave accents (`...`) is taken as the command to be executed and is replaced with the output from the command. The command is written using the usual quoting conventions except that a ` must be escaped using a \. For example,

ls `echo "$1"`

is equivalent to

ls $1

Command substitution occurs in all contexts where parameter substitution occurs (including *here* documents) and the treatment of the resulting text is the same in both cases. This mechanism allows string processing commands to be used within shell procedures. An example of such a command is *basename* which removes a specified suffix from a string. For example,

basename main.c .c

will print the string *main*. Its use is illustrated by the following fragment from a *cc* command.

case $A in
   ...
   *.c)   B=`basename $A .c`
   ...
esac

that sets **B** to the part of $A with the suffix .c stripped.

Here are some composite examples.

- **for i in `ls —t`; do ...**
  The variable **i** is set to the names of files in time order, most recent first.

- **set `date`; echo $6 $2 $3, $4**
  will print, e.g., *1977 Nov 1, 23:59:59*

3.4 Evaluation and quoting

The shell is a macro processor that provides parameter substitution, command substitution and file name generation for the arguments to commands. This section discusses the order in which these evaluations occur and the effects of the various quoting mechanisms.

Commands are parsed initially according to the grammar given in appendix A. Before a command is executed the following substitutions occur.

- parameter substitution, e.g. **$user**
- command substitution, e.g. **`pwd`**
  Only one evaluation occurs so that if, for example, the value of the variable **X** is the string *$y* then

    echo $X

  will echo *$y*.

- blank interpretation
  Following the above substitutions the resulting characters are broken into non-blank words (*blank interpretation*). For this purpose 'blanks' are the characters of the string **$IFS**. By default, this string consists of blank, tab and newline. The null string is not regarded as a word unless it is quoted. For example,

    echo ''

  will pass on the null string as the first argument to *echo*, whereas

    echo $null

  will call *echo* with no arguments if the variable **null** is not set or set to the null string.

- file name generation
  Each word is then scanned for the file pattern characters *****, **?** and **[...]** and an alphabetical list of file names is generated to replace the word. Each such file name is a separate argument.

The evaluations just described also occur in the list of words associated with a **for** loop. Only substitution occurs in the *word* used for a **case** branch.

As well as the quoting mechanisms described earlier using \ and `'...'` a third quoting mechanism is provided using double quotes. Within double quotes parameter and command substitution occurs but file name generation and the interpretation of blanks does not. The following characters have a special meaning within double quotes and may be quoted using \.

$	parameter substitution
`	command substitution
"	ends the quoted string
\	quotes the special characters $ ` " \

For example,

    echo "$x"

will pass the value of the variable x as a single argument to *echo*. Similarly,

echo "$*"

will pass the positional parameters as a single argument and is equivalent to

echo "$1 $2 ..."

The notation $@ is the same as $* except when it is quoted.

echo "$@"

will pass the positional parameters, unevaluated, to *echo* and is equivalent to

echo "$1" "$2" ...

The following table gives, for each quoting mechanism, the shell metacharacters that are evaluated.

			metacharacter			
	\	$	*	`	"	'
'	n	n	n	n	n	t
`	y	n	n	t	n	n
"	y	y	n	y	t	n

t	terminator
y	interpreted
n	not interpreted

Figure 2. Quoting mechanisms

In cases where more than one evaluation of a string is required the built-in command *eval* may be used. For example, if the variable X has the value $y, and if y has the value *pqr* then

eval echo $X

will echo the string *pqr*.

In general the *eval* command evaluates its arguments (as do all commands) and treats the result as input to the shell. The input is read and the resulting command(s) executed. For example,

wg='eval who I grep'
$wg fred

is equivalent to

who I grep fred

In this example, *eval* is required since there is no interpretation of metacharacters, such as I , **following substitution.**

3.5 Error handling

The treatment of errors detected by the shell depends on the type of error and on whether the shell is being used interactively. An interactive shell is one whose input and output are connected to a terminal (as determined by *gtty* (2)). A shell invoked with the −i flag is also interactive.

Execution of a command (see also 3.7) may fail for any of the following reasons.

● Input output redirection may fail. For example, if a file does not exist or cannot be created.

- The command itself does not exist or cannot be executed.
- · The command terminates abnormally, for example, with a "bus error" or "memory fault". See Figure 2 below for a complete list of UNIX signals.
- The command terminates normally but returns a non-zero exit status.

In all of these cases the shell will go on to execute the next command. Except for the last case an error message will be printed by the shell. All remaining errors cause the shell to exit from a command procedure. An interactive shell will return to read another command from the terminal. Such errors include the following.

- Syntax errors. e.g., if ... then ... done
- A signal such as interrupt. The shell waits for the current command, if any, to finish execution and then either exits or returns to the terminal.
- Failure of any of the built-in commands such as *cd.*

The shell flag —e causes the shell to terminate if any error is detected.

1	hangup
2	interrupt
3*	quit
4*	illegal instruction
5*	trace trap
6*	IOT instruction
7*	EMT instruction
8*	floating point exception
9	kill (cannot be caught or ignored)
10*	bus error
11*	segmentation violation
12*	bad argument to system call
13	write on a pipe with no one to read it
14	alarm clock
15	software termination (from *kill* (1))

Figure 3. UNIX signals

Those signals marked with an asterisk produce a core dump if not caught. However, the shell itself ignores quit which is the only external signal that can cause a dump. The signals in this list of potential interest to shell programs are 1, 2, 3, 14 and 15.

3.6 Fault handling

Shell procedures normally terminate when an interrupt is received from the terminal. The *trap* command is used if some cleaning up is required, such as removing temporary files. For example,

        trap 'rm /tmp/ps$$; exit' 2

sets a trap for signal 2 (terminal interrupt), and if this signal is received will execute the commands

        rm /tmp/ps$$; exit

exit is another built-in command that terminates execution of a shell procedure. The *exit* is required; otherwise, after the trap has been taken, the shell will resume executing the procedure at the place where it was interrupted.

UNIX signals can be handled in one of three ways. They can be ignored, in which case the signal is never sent to the process. They can be caught, in which case the process must decide what action to take when the signal is received. Lastly, they can be left to cause termination of

the process without it having to take any further action. If a signal is being ignored on entry to the shell procedure, for example, by invoking it in the background (see 3.7) then *trap* commands (and the signal) are ignored.

The use of *trap* is illustrated by this modified version of the *touch* command (Figure 4). The cleanup action is to remove the file **junk$$**.

```
flag =
trap 'rm −f junk$$; exit' 1 2 3 15
for i
do case $i in
    −c)  flag=N ;;
    *)   if test −f $i
         then    ln $i junk$$; rm junk$$
         elif test $flag
         then    echo file \'$i\' does not exist
         else    >$i
         fi
    esac
done
```

Figure 4. The touch command

The *trap* command appears before the creation of the temporary file; otherwise it would be possible for the process to die without removing the file.

Since there is no signal 0 in UNIX it is used by the shell to indicate the commands to be executed on exit from the shell procedure.

A procedure may, itself, elect to ignore signals by specifying the null string as the argument to trap. The following fragment is taken from the *nohup* command.

```
trap '' 1 2 3 15
```

which causes *hangup, interrupt, quit* and *kill* to be ignored both by the procedure and by invoked commands.

Traps may be reset by saying

```
trap 2 3
```

which resets the traps for signals 2 and 3 to their default values. A list of the current values of traps may be obtained by writing

```
trap
```

The procedure *scan* (Figure 5) is an example of the use of *trap* where there is no exit in the trap command. *scan* takes each directory in the current directory, prompts with its name, and then executes commands typed at the terminal until an end of file or an interrupt is received. Interrupts are ignored while executing the requested commands but cause termination when *scan* is waiting for input.

```
d=`pwd`
for i in *
do if test −d $d/$i
    then cd $d/$i
         while echo "$i:"
                 trap exit 2
                 read x
         do trap : 2; eval $x; done
   fi
done
```

Figure 5. The scan command

read x is a built-in command that reads one line from the standard input and places the result in the variable **x**. It returns a non-zero exit status if either an end-of-file is read or an interrupt is received.

3.7 Command execution

To run a command (other than a built-in) the shell first creates a new process using the system call *fork*. The execution environment for the command includes input, output and the states of signals, and is established in the child process before the command is executed. The built-in command *exec* is used in the rare cases when no fork is required and simply replaces the shell with a new command. For example, a simple version of the *nohup* command looks like

```
trap " 1 2 3 15
exec $*
```

The *trap* turns off the signals specified so that they are ignored by subsequently created commands and *exec* replaces the shell by the command specified.

Most forms of input output redirection have already been described. In the following *word* is only subject to parameter and command substitution. No file name generation or blank interpretation takes place so that, for example,

```
echo ... >*.c
```

will write its output into a file whose name is ***.c**. Input output specifications are evaluated left to right as they appear in the command.

> *word*   The standard output (file descriptor 1) is sent to the file *word* which is created if it does not already exist.

>> *word*   The standard output is sent to file *word*. If the file exists then output is appended (by seeking to the end); otherwise the file is created.

< *word*   The standard input (file descriptor 0) is taken from the file *word*.

<< *word*   The standard input is taken from the lines of shell input that follow up to but not including a line consisting only of *word*. If *word* is quoted then no interpretation of the document occurs. If *word* is not quoted then parameter and command substitution occur and \ is used to quote the characters \ $ ` and the first character of *word*. In the latter case **newline** is ignored (c.f. quoted strings).

>& *digit*   The file descriptor *digit* is duplicated using the system call *dup* (2) and the result is used as the standard output.

<& *digit*   The standard input is duplicated from file descriptor *digit*.

<&-	The standard input is closed.
>&-	The standard output is closed.

Any of the above may be preceded by a digit in which case the file descriptor created is that specified by the digit instead of the default 0 or 1. For example,

> ... 2>file

runs a command with message output (file descriptor 2) directed to *file*.

> ... 2>&1

runs a command with its standard output and message output merged. (Strictly speaking file descriptor 2 is created by duplicating file descriptor 1 but the effect is usually to merge the two streams.)

The environment for a command run in the background such as

> list *.c | lpr &

is modified in two ways. Firstly, the default standard input for such a command is the empty file **/dev/null**. This prevents two processes (the shell and the command), which are running in parallel, from trying to read the same input. Chaos would ensue if this were not the case. For example,

> ed file &

would allow both the editor and the shell to read from the same input at the same time.

The other modification to the environment of a background command is to turn off the QUIT and INTERRUPT signals so that they are ignored by the command. This allows these signals to be used at the terminal without causing background commands to terminate. For this reason the UNIX convention for a signal is that if it is set to 1 (ignored) then it is never changed even for a short time. Note that the shell command *trap* has no effect for an ignored signal.

3.8 Invoking the shell

The following flags are interpreted by the shell when it is invoked. If the first character of argument zero is a minus, then commands are read from the file **.profile**.

−c *string*
> If the −c flag is present then commands are read from *string*.

−s
> If the −s flag is present or if no arguments remain then commands are read from the standard input. Shell output is written to file descriptor 2.

−i
> If the −i flag is present or if the shell input and output are attached to a terminal (as told by *gtty*) then this shell is *interactive*. In this case TERMINATE is ignored (so that **kill 0** does not kill an interactive shell) and INTERRUPT is caught and ignored (so that **wait** is interruptable). In all cases QUIT is ignored by the shell.

Acknowledgements

References

1. B. W. Kernighan, *UNIX for Beginners*, Bell Laboratories internal memorandum (1978).

2. K. Thompson and D. M. Ritchie, *UNIX Programmer's Manual*, Bell Laboratories (1978). Seventh Edition.

3. K. Thompson, "The UNIX Command Language," pp. 375-384 in *Structured Programming—Infotech State of the Art Report*, Infotech International Ltd., Nicholson House, Maidenhead, Berkshire, England (March 1975).

4. J. R. Mashey, *PWB/UNIX Shell Tutorial*, Bell Laboratories internal memorandum (September 30, 1977).

5. D. F. Hartley (Ed.), *The Cambridge Multiple Access System — Users Reference Manual*, University Mathematical Laboratory, Cambridge, England (1968).

6. P. A. Crisman (Ed.), *The Compatible Time-Sharing System*, M.I.T. Press, Cambridge, Mass. (1965).

Appendix A - Grammar

item:	*word*
	input-output
	name = *value*
simple-command:	*item*
	simple-command item
command:	*simple-command*
	(*command-list*)
	{ *command-list* }
	for *name* **do** *command-list* **done**
	for *name* **in** *word* ... **do** *command-list* **done**
	while *command-list* **do** *command-list* **done**
	until *command-list* **do** *command-list* **done**
	case *word* **in** *case-part* ... **esac**
	if *command-list* **then** *command-list else-part* **fi**
pipeline:	*command*
	pipeline	*command*
andor:	*pipeline*
	andor && *pipeline*
	andor || *pipeline*
command-list:	*andor*
	command-list ;
	command-list &
	command-list ; *andor*
	command-list & *andor*
input-output:	> *file*
	< *file*
	>> *word*
	<< *word*
file:	*word*
	& *digit*
	& −
case-part:	*pattern*) *command-list* ;;
pattern:	*word*
	pattern	*word*
else-part:	**elif** *command-list* **then** *command-list else-part*
	else *command-list*
	empty
empty:	
word:	a sequence of non-blank characters
name:	a sequence of letters, digits or underscores starting with a letter
digit:	**0 1 2 3 4 5 6 7 8 9**

Appendix B - Meta-characters and Reserved Words

a) syntactic

|	pipe symbol
&&	'andf' symbol
||	'orf' symbol
;	command separator
;;	case delimiter
&	background commands
( )	command grouping
<	input redirection
<<	input from a here document
>	output creation
>>	output append

b) patterns

*	match any character(s) including none
?	match any single character
[...]	match any of the enclosed characters

c) substitution

${...}	substitute shell variable
`...`	substitute command output

d) quoting

\	quote the next character
'...'	quote the enclosed characters except for '
"..."	quote the enclosed characters except for $ ` \ "

e) reserved words

if then else elif fi
case in esac
for while until do done
{ }

ADDENDUM

Since its release there have been several changes to the shell; most of these changes were minor bug fixes and should be relatively transparent (painless) to the user. However, a few of these enhancements are significant and are described below.

A. When invoked, the shell now searches the environment for the variable SHELL. If SHELL is found, and has an ``r'' in the simple part of its value (the part after the last ``/'') then the shell becomes restricted. Restricted logins, such as games, should include the following in their .profile files:

        SHELL=/usr/rsh
        export SHELL   .

This will cause any unprotected escapes from games and the like to get a restricted shell when exec'd.

B. A new option (--) has been added to the set command. This option indicates that the flags are to remain unchanged and is useful in setting $1 to a string beginning with a minus (-).

C. File name generation has been enhanced to include a not operator. If a ``!'' immediately follows the left of a pair of square brackets ([]) then any character not enclosed within the brackets is matched. For example,

        ls *[!o]

will list all files not ending in `o'.

D. At some point, some rather strange bugs were introduced in both continue n and break n. These problems have been fixed and both commands now behave exactly as documented. In particular loops such as:

        while echo hi
        do
                continue
        done
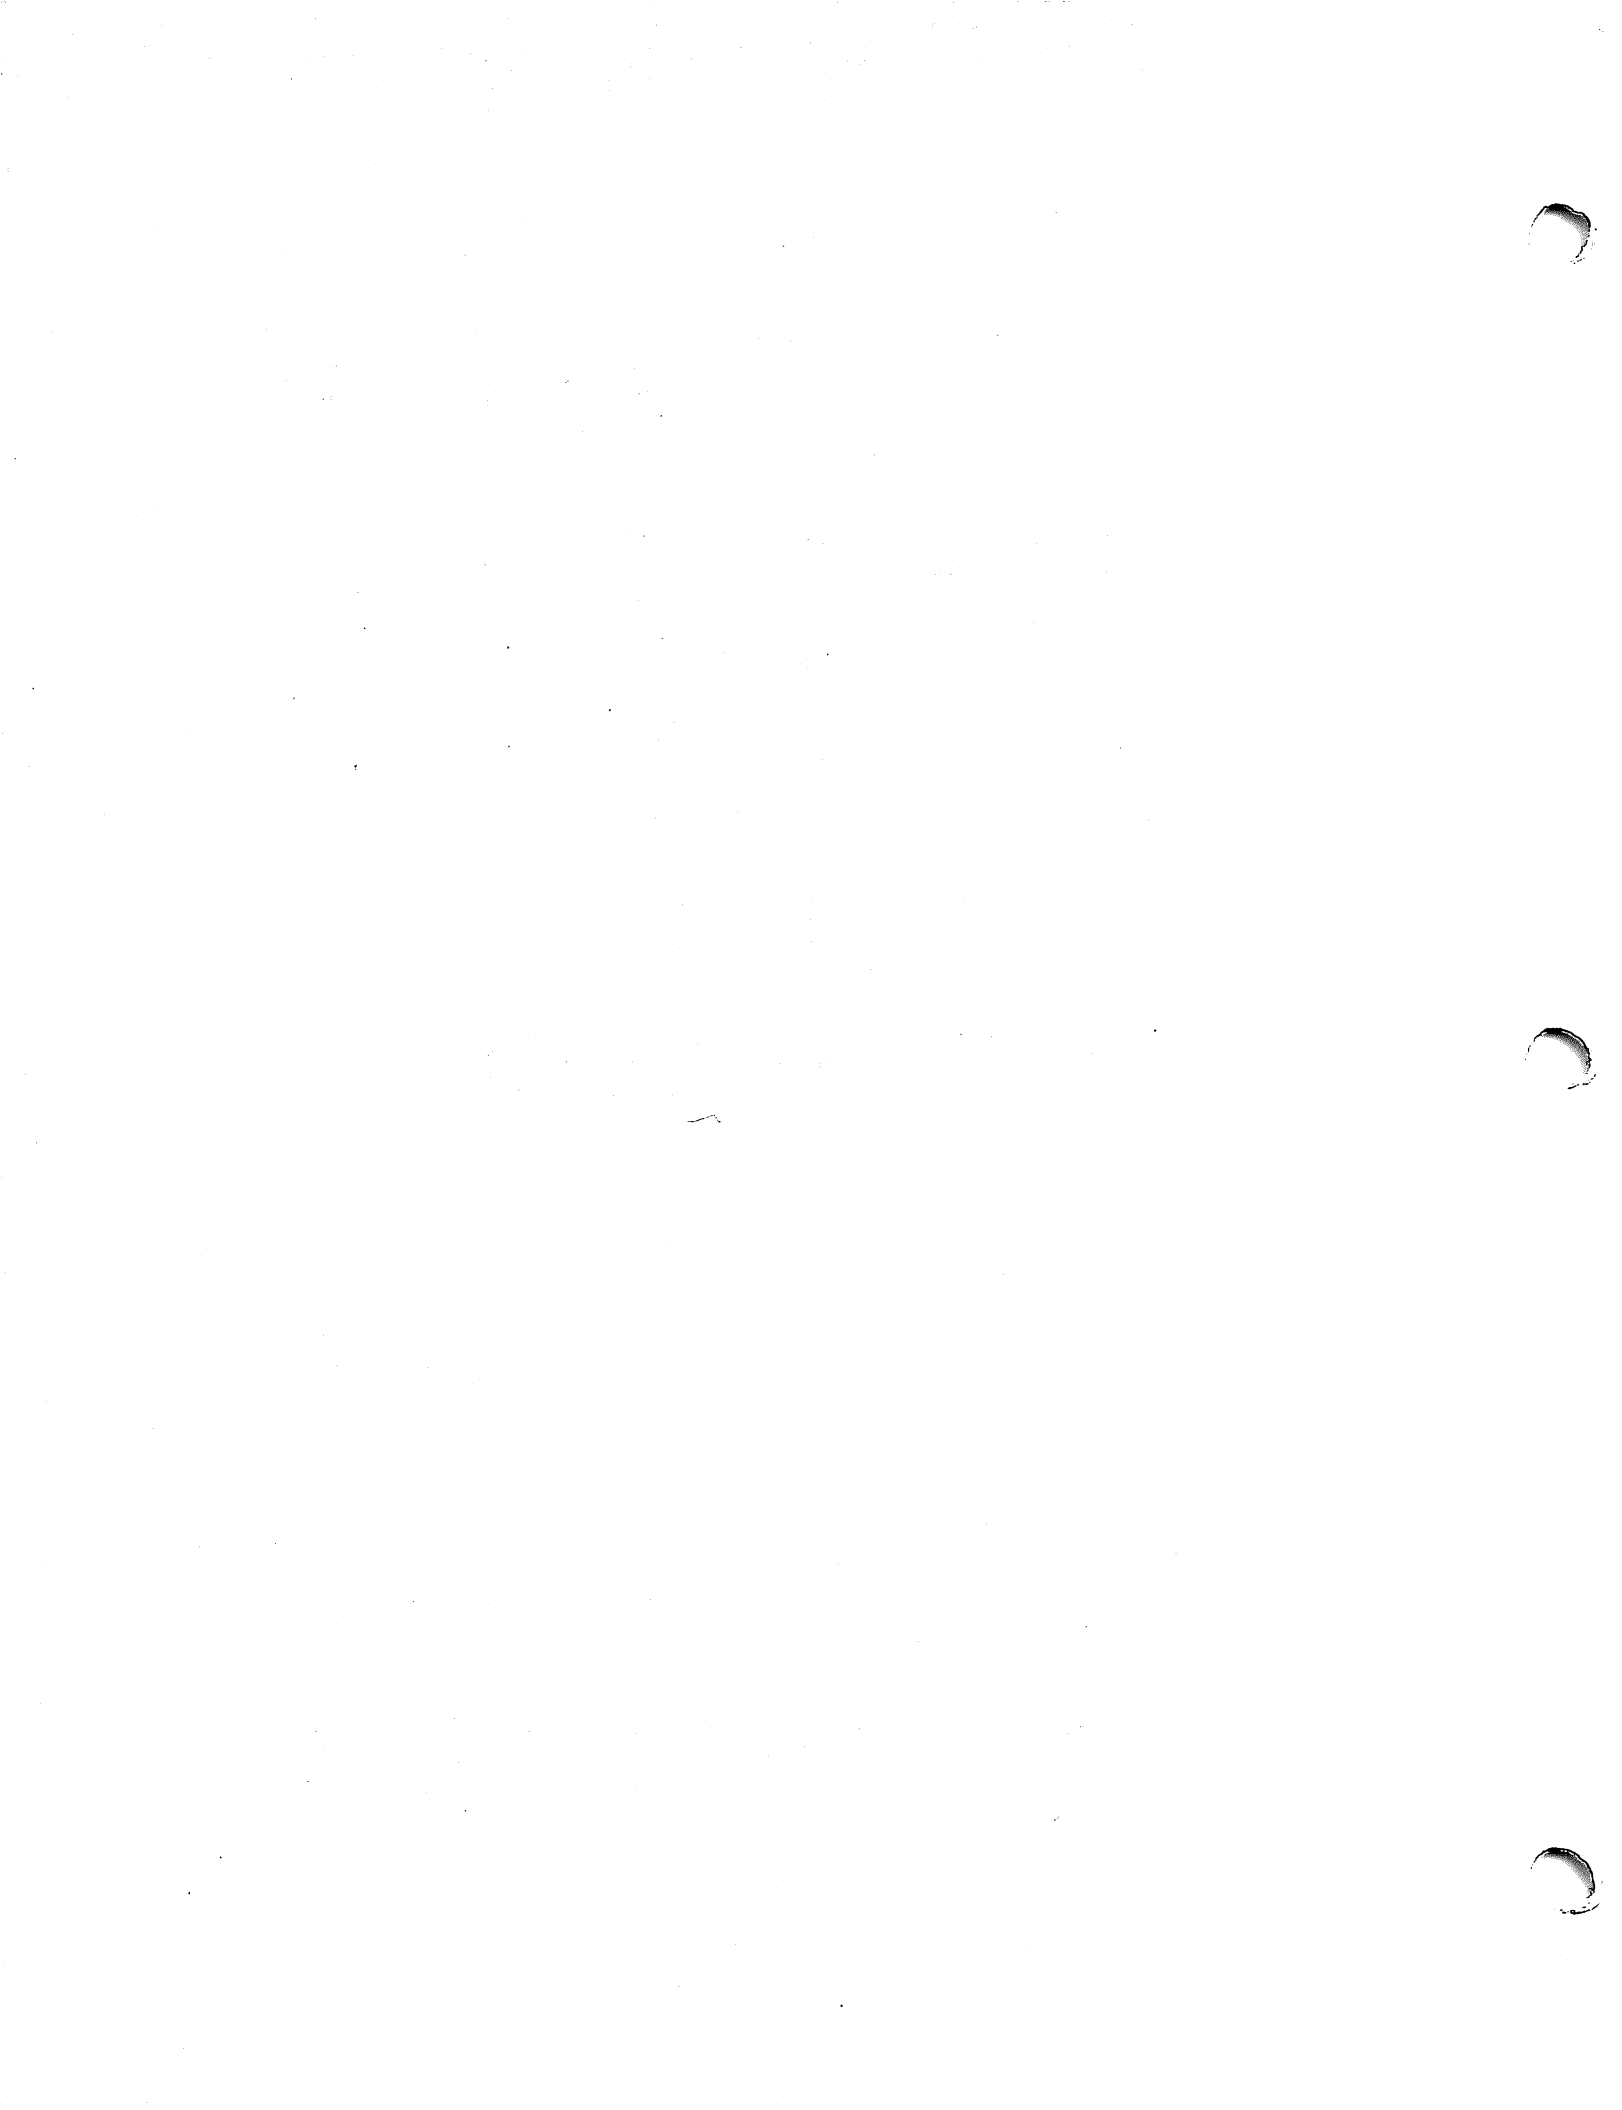
will echo an infinite number of hi's, and interrupting the loop will no longer cause the shell to behave as if set -n had been typed.

E. The shell will once again execute the commands contained in both /etc/profile and in .profile whenever invoked with an argument zero ($0) containing a ``-'' as the first character; thus commands like

        su - name

will once again behave as documented.

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 System[1] 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
$-$o*list*	Print only pages whose page numbers appear in *list*, 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.
$-$n*N*	Number first generated page N.
$-$s*N*	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.
$-$m*name*	Prepends the macro file /usr/lib/tmac.*name* to the input *files*.
$-$ra*N*	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

−T *name* 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.

−p *N* Print all characters in point size *N* while retaining all prescribed spacings and motions, to reduce phototypesetter elasped 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 EQN[2] (for NROFF and TROFF respectively), and the table-construction preprocessor TBL[3]. A reverse-line postprocessor COL[4] 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. TK[4] is a 37 Teletype simulator postprocessor for printing NROFF output on a Tektronix 4014. TCAT[4] 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. GCAT[4] 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].

<div align="right">Joseph F. Ossanna</div>

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 $\pm N$	10 point	previous	E	Point size; also \s$\pm N$.†
.ss N	12/36 em	ignored	E	Space-character size set to N/36 em.†
.cs $F N M$	off	-	P	Constant character space (width) mode (font F).†
.bd $F N$	off	-	P	Embolden font F by $N-1$ units.†
.bd S $F N$	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 \f$x$,\f(xx,\f$N$.
.fp $N F$	R,I,B,S	ignored	-	Font named F mounted on physical position $1 \leqslant N \leqslant 4$.
3. Page Control				
.pl $\pm N$	11 in	11 in	v	Page length.
.bp $\pm N$	N=1	-	B‡,v	Eject current page; next page number N.
.pn $\pm N$	N=1	ignored	-	Next page number N.
.po $\pm N$	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 $\pm N$	none	internal	D,v	Return (upward only) 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 $\pm N$	6.5 in	previous	E,m	Line length.
.in $\pm N$	N=0	previous	B,E,m	Indent.
.ti $\pm N$	-	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 *xx*	-	ignored	-	Remove request, macro, or string.
.rn *xx yy*	-	ignored	-	Rename request, macro, or string *xx* to *yy*.
.di *xx*	-	end	D	Divert output to macro *xx*.
.da *xx*	-	end	D	Divert and append to *xx*.
.wh *N xx*	-	-	v	Set location trap; negative is w.r.t. page bottom.
.ch *xx N*	-	-	v	Change trap location.
.dt *N xx*	-	off	D,v	Set a diversion trap.
.it *N xx*	-	off	E	Set an input-line count trap.
.em *xx*	none	none	-	End macro is *xx*.

8. Number Registers

.nr *R ±N M*	-	-	u	Define and set number register *R*; auto-increment by *M*.
.af *R c*	arabic	-	-	Assign format to register *R* (*c*=1, i, I, a, A).
.rr *R*	-	-	-	Remove register *R*.

9. Tabs, Leaders, and Fields

.ta *Nt* ...	0.8; 0.5in	none	E,m	Tab settings; *left* type, unless *t*=R(right), C(centered).
.tc *c*	none	none	E	Tab repetition character.
.lc *c*	.	none	E	Leader repetition character.
.fc *a b*	off	off	-	Set field delimiter *a* and pad character *b*.

10. Input and Output Conventions and Character Translations

.ec *c*	\	\	-	Set escape character.
.eo	on	-	-	Turn off escape character mechanism.
.lg *N*	-; on	on	-	Ligature mode on if *N*>0.
.ul *N*	off	*N*=1	E	Underline (italicize in TROFF) *N* input lines.
.cu *N*	off	*N*=1	E	Continuous underline in NROFF; like ul in TROFF.
.uf *F*	Italic	Italic	-	Underline font set to *F* (to be switched to by ul).
.cc *c*	.	.	E	Set control character to *c*.
.c2 *c*	'	'	E	Set nobreak control character to *c*.
.tr *abcd*....	none	-	O	Translate *a* to *b*, 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 *N*	hyphenate	hyphenate	E	Hyphenate; *N* = mode.
.hc *c*	\%	\%	E	Hyphenation indicator character *c*.
.hw *word1* ...		ignored	-	Exception words.

14. Three Part Titles.

.tl '*left*'*center*'*right*'	-	-	-	Three part title.
.pc *c*	%	off	-	Page number character.
.lt ±*N*	6.5 in	previous	E,m	Length of title.

15. Output Line Numbering.

.nm ±*N M S I*	off	E		Number mode on or off, set parameters.
.nn *N*	-	*N*=1	E	Do not number next *N* lines.

16. Conditional Acceptance of Input

.if *c anything*	-	-		If condition *c* true, accept *anything* as input, for multi-line use \\{*anything*\\}.

Request Form	Initial Value	If No Argument	Notes	Explanation
.if !c *anything*	-	-		If condition *c* false, accept *anything*.
.if *N anything*	-	u		If expression $N > 0$, accept *anything*.
.if !*N anything*	-	u		If expression $N \leqslant 0$, accept *anything*.
.if '*string1*'*string2*' *anything*	-			If *string1* identical to *string2*, accept *anything*.
.if !'*string1*'*string2*' *anything*	-			If *string1* not identical to *string2*, accept *anything*.
.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*	*N*=0	previous	-	Environment switched (*push down*).

18. Insertions from the Standard Input

.rd *prompt*	-	*prompt*=BEL	-	Read insertion.
.ex	-	-	-	Exit from NROFF/TROFF.

19. Input/Output File Switching

.so *filename*	-	-		Switch source file (*push down*).
.nx *filename*	end-of-file	-		Next file.
.pi *program*	-	-		Pipe output to *program* (NROFF only).

20. Miscellaneous

.mc *c N*	-	off	E,m	Set margin character *c* and separation *N*.
.tm *string*	-	newline	-	Print *string* on terminal (UNIX standard message output).
.ig *yy*	-	.*yy*=..	-	Ignore till call of *yy*.
.pm *t*	-	all	-	Print macro names and sizes; if *t* 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 *current* escape character.
2.1	\´	´ (acute accent); equivalent to \(aa
2.1	\`	` (grave accent); equivalent to \(ga
2.1	\-	− Minus sign in the *current* 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 $1 \leqslant N \leqslant 9$
13	\%	Default optional hyphenation character
2.1	\(xx	Character named xx
7.1	*x, *(xx	Interpolate string x or xx
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	\f$x$,\f(xx,\f$N$	Change to font named x or xx, or position N
11.1	\h'N'	Local horizontal motion; move right N *(negative left)*
11.3	\k$x$	Mark horizontal *input* place in register x
12.4	\l'Nc'	Horizontal line drawing function (optionally with c)
12.4	\L'Nc'	Vertical line drawing function (optionally with c)
8	\n$x$,\n(xx	Interpolate number register x or xx
12.1	\o'abc...'	Overstrike characters $a, b, c, ...$
4.1	\p	Break and spread output line
11.1	\r	Reverse 1 em vertical motion (reverse line in NROFF)
2.3	\s$N$,\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 N *(negative up)*
11.2	\w'string'	Interpolate width of *string*
5.2	\x'N'	Extra line-space function *(negative before, positive after)*
12.2	\z$c$	Print c with zero width (without spacing)
16	\{	Begin conditional input
16	\}	End conditional input
10.7	\(newline)	Concealed (ignored) newline
-	\X	X, any character *not* listed above

The escape sequences \\, \., \", \$, *, \a, \n, \t, and \(newline) are interpreted in *copy mode* (§7.2).

Predefined General Number Registers

Section Reference	Register Name	Description
3	%	Current page number.
11.2	ct	Character type (set by *width* 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 *input* 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 *width* function).
11.2	st	Height of string above base line (generated by *width* function).
-	yr	Last two digits of current year.

Predefined Read-Only Number Registers

Section Reference	Register Name	Description
7.3	.$	Number of arguments available at the current macro level.
-	.A	Set to 1 in TROFF, if −a option used; always 1 in NROFF.
11.1	.H	Available horizontal resolution in basic units.
-	.T	Set to 1 in NROFF, if −T 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 $\backslash x'N'$.
-	.c	Number of *lines* read from current input file.
7.4	.d	Current vertical place in current diversion; equal to **nl**, 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 \nx, 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 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. ±*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 @, #, ", ´, `, <, >, \, {, }, ˜, ˆ, 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 \f*x,* \f(*xx,* or \f*N* 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 \s*N* at the desired point to set the size to *N,* or a \s±*N* $(1 \leqslant N \leqslant 9)$ 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.

Request Form	*Initial Value*	*If No Argument*	*Notes**	*Explanation*
.ps ±*N*	10 point	previous	E	Point size set to ±*N.* Alternatively imbed \s*N* or \s±*N.* 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.
.ss *N*	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.
.cs *F N M*	off	-	P	Constant character space (width) mode is set on for font *F* (if mounted); the width of every character will be taken to be *N*/36 ems. If *M* is absent, the em is that of the character's point size; if *M* is given, the em is *M*-points. All affected characters are centered in this space, including those with an actual width larger than this space. Special Font characters occurring while the current font is *F* are also so treated. If *N* 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.
.bd *F N*	off	-	P	The characters in font *F* will be artificially emboldened by printing each one twice, separated by *N*−1 basic units. A reasonable value for *N* is 3 when the character size is in the vicinity of 10 points. If *N* is missing the embolden mode is turned off. The column heads above were printed with **.bd I 3.** 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 S *F N*	off	-	P	The characters in the Special Font will be emboldened whenever the current font is *F*. This manual was printed with **.bd S B 3**. The mode must be still or again in effect when the characters are physically printed.
.ft *F*	Roman	previous	E	Font changed to *F*. Alternatively, imbed \f*F*. The font name **P** is reserved to mean the previous font.
.fp *N F*	R,I,B,S	ignored	-	Font position. This is a statement that a font named *F* is mounted on position *N* (1-4). It is a fatal error if *F* 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 ±*N*	11 in	11 in	v	Page length set to ±*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 ±*N*	*N*=1	-	B*,v	Begin page. The current page is ejected and a new page is begun. If ±*N* is given, the new page number will be ±*N*. Also see request **ns**.
.pn ±*N*	*N*=1	ignored	-	Page number. The next page (when it occurs) will have the page number ±*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 ±*N*	0; 26/27 in†	previous	v	Page offset. The current *left margin* is set to ±*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 *N*	-	*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.

.mk R	none	internal	D	Mark the *current* vertical place in an internal register (both associated with the current diversion level), or in register R, if given. See rt request.
.rt $\pm N$	none	internal	D,v	Return *upward only* to a marked vertical place in the current diversion. If $\pm N$ (w.r.t. current place) is given, the place is $\pm N$ from the top of the page or diversion or, if N is absent, to a place marked by a previous **mk**. Note that the **sp** request (§5.3) may be used in all cases instead of **rt** by spacing to the absolute place stored in a explicit register; e. g. using the sequence **.mk** R ... **.sp** $|\backslash nRu$.

4. Text Filling, Adjusting, and Centering

4.1. Filling and adjusting. Normally, words are collected from input text lines and assembled into a output text line until some word doesn't fit. An attempt is then made the 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 a 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.

Request Form	Initial Value	If No Argument	Notes	Explanation
.br	-	-	B	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 *neither* filled *nor* adjusted. Input text lines are copied directly to output lines *without regard* 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 *c* 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 *N* input text lines within the current (line-length minus indent). If *N*=0, any residual count is cleared. A break occurs after each of the *N* 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 *V*. Transient *extra* vertical space available with \x'*N*' (see above).
.ls *N*	*N*=1	previous	E	*Line* spacing set to ±*N*. *N*−1 *V*s *(blank lines)* 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 *N*	-	*N*=1 *V*	B,v	Space vertically in *either* direction. If *N* is negative, the motion is *backward* (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 *N*	-	*N*=1 *V*	v	Save a contiguous vertical block of size *N*. If the distance to the next trap is greater than *N*, *N* vertical space is output. No-space mode has *no* effect. If this distance is less than *N*, no vertical space is immediately output, but *N* is remembered for later output (see **os**). Subsequent **sv** requests will overwrite any still remembered *N*.
.os	-	-	-	Output saved vertical space. No-space mode has *no* 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 *without* 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 ±*N*	6.5 in	previous	E,m	Line length is set to ±*N*. In TROFF the maximum (line-length)+(page-offset) is about 7.54 inches.
.in ±*N*	*N*=0	previous	B,E,m	Indent is set to ±*N*. The indent is prepended to each output line.
.ti ±*N*	-	ignored	B,E,m	Temporary indent. The *next* output text line will be indented a distance ±*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 Nth argument ($1 \leqslant N \leqslant 9$). 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.

Request Form	Initial Value	If No Argument	Notes	Explanation
.de *xx yy*	-	*.yy*=..	-	Define or redefine the macro *xx*. The contents of the macro begin on the next input line. Input lines are copied in *copy mode* until the definition is terminated by a line beginning with *.yy*, whereupon the macro *yy* is called. In the absence of *yy*, the definition is terminated by a line beginning with "..". A macro may contain **de** requests provided the terminating macros differ or the contained definition terminator is concealed. ".." can be concealed as \\.. which will copy as \.. and be reread as "..".
.am *xx yy*	-	*.yy*=..	-	Append to macro (append version of **de**).
.ds *xx string*	-	ignored	-	Define a string *xx* containing *string*. Any initial double-quote in *string* is stripped off to permit initial blanks.
.as *xx string*	-	ignored	-	Append *string* to string *xx* (append version of **ds**).
.rm *xx*	-	ignored	-	Remove request, macro, or string. The name *xx* is removed from the name list and any related storage space is freed. Subsequent references will have no effect.
.rn *xx yy*	-	ignored	-	Rename request, macro, or string *xx* to *yy*. If *yy* exists, it is first removed.
.di *xx*	-	end	D	Divert output to macro *xx*. 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 *xx* | - | end | D | Divert, appending to *xx* (append version of **di**). |

| .wh *N xx* | - | - | v | Install a trap to invoke *xx* at page position *N;* a *negative N* will be interpreted with respect to the page *bottom*. Any macro previously planted at *N* is replaced by *xx*. A zero *N* refers to the *top* of a page. In the absence of *xx*, the first found trap at *N*, if any, is removed. |

| .ch *xx N* | - | - | v | Change the trap position for macro *xx* to be *N*. In the absence of *N*, the trap, if any, is removed. |

| .dt *N xx* | - | off | D,v | Install a diversion trap at position *N* in the *current* diversion to invoke macro *xx*. Another **dt** will redefine the diversion trap. If no arguments are given, the diversion trap is removed. |

| .it *N xx* | - | off | E | Set an input-line-count trap to invoke the macro *xx* after *N* lines of *text* 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 *xx* | none | none | - | The macro *xx* will be invoked when all input has ended. The effect is the same as if the contents of *xx* 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
\n*x*	none	N
\n(*xx*	none	N
\n+*x*	*x* incremented by M	$N+M$
\n−*x*	*x* decremented by M	$N-M$
\n+(*xx*	*xx* incremented by M	$N+M$
\n−(*xx*	*xx* decremented by M	$N-M$

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**.

Request Form	*Initial Value*	*If No Argument*	*Notes*	*Explanation*
.nr *R* ±*N M*	-		u	The number register *R* is assigned the value ±*N* with respect to the previous value, if any. The increment for auto-incrementing is set to *M*.

.af *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 *next-string*
Left	D	Following D
Right	$D-W$	Right adjusted within D
Centered	$D-W/2$	Centered on right end of D

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.

Request Form	Initial Value	If No Argument	Notes	Explanation
.ta Nt ...	0.8; 0.5in	none	E,m	Set tab stops and types. $t=R$, right adjusting; $t=C$, centering; t absent, ieft 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 c	none	none	E	The tab repetition character becomes c, or is removed specifying motion.
.lc c	.	none	E	The leader repetition character becomes c, or is removed specifying motion.
.fc a b	off	off	-	The field delimiter is set to a; the padding indicator is set to the *space* character or to b, 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**.

Request Form	Initial Value	If No Argument	Notes	Explanation
.ec c	\	\	-	Set escape character to \, or to c, 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, \(Fi, and \(Fl respectively. The ligature mode is normally on in TROFF, and *automatically* invokes ligatures during input.

Request Form	Initial Value	If No Argument	Notes	Explanation
.lg N	off; on	on	-	Ligature mode is turned on if N is absent or non-zero, and turned off if $N=0$. If $N=2$, only the two-character ligatures are automatically invoked. Ligature mode is inhibited for request, macro, string, register, or file names, and in *copy mode*. 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 \f$F$, the underline font may be selected by **ul** and **cu**. Underlining is restricted to an output-device-dependent subset of *reasonable* characters.

Request Form	Initial Value	If No Argument	Notes	Explanation
.ul *N*	off	*N*=1	E	Underline in NROFF (italicize in TROFF) the next *N* input text lines. Actually, switch to *underline* font, saving the current font for later restoration; *other* 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 *not* decrement *N*. If *N*>1, there is the risk that a trap interpolated macro may provide text lines within the span; environment switching can prevent this.
.cu *N*	off	*N*=1	E	A variant of **ul** that causes *every* character to be underlined in NROFF. Identical to **ul** in TROFF.
.uf *F*	Italic	Italic	-	Underline font set to *F.* In NROFF, *F* may *not* 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.

Request Form	Initial Value	If No Argument	Notes	Explanation
.cc *c*	.	.	E	The basic control character is set to *c*, or reset to ".".
.c2 *c*	'	'	E	The *nobreak* control character is set to *c*, 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).

Request Form	Initial Value	If No Argument	Notes	Explanation
.tr *abcd....*	none	-	O	Translate *a* into *b*, *c* into *d*, 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 *input* to *output* 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 TROFF	NROFF	Horizontal Local Motion	Effect in TROFF	NROFF
\v'N'	Move distance N		\h'N' \(space) \0	Move distance N Unpaddable space-size space Digit-size space	
\u \d \r	½ em up ½ em down 1 em up	½ line up ½ line down 1 line up	| \\^	1/6 em space 1/12 em space	ignored ignored

As an example, E^2 could be generated by the sequence E\s−2\v'−0.4m'2\v'0.4m'\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 \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, .ti −\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 \n(stu−\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 \k*x* will cause the current horizontal position in the *input line* to be stored in register *x*. As an example, the construction \kx*word*\h'|\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 *over-strike* function \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, \o'e\"' produces é, and \o'\(mo\(sl' produces ∉.

12.2. Zero-width characters. The function \z*c* will output *c* without spacing over it, and can be used to produce left-aligned overstruck combinations. As examples, \z\(ci\(pl will produce ⊕, and \(br\z\(rn\(ul\(br will produce the smallest possible constructed box □.

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 \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 (½ line in NROFF). For example, \b'\(lc\(lf'E\|\b'\(rc\(rf'\x'−0.5m'\x'0.5m' produces ⎡E⎤.

12.4. Line drawing. The function \l'Nc' will draw a string of repeated *c*'s towards the right for a distance *N*. (\l is \(lower case L). If *c* looks like a continuation of an expression for *N*, it may insulated from *N* with a \&. 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*/(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
\\$1\l'|0\(ul'
..
```

or one to draw a box around a string

```
.de bx
\(br\|\\$1\|\(br\ l ´|0\(rn´\ l ´|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 ½-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´\l´\\n(.lu+1n\(ul´\L´−|\\nau+1´\l´|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.

Request Form	Initial Value	If No Argument	Notes	Explanation
.nh	hyphenate	-	E	Automatic hyphenation is turned off.
.hy N	on, N=1	on, N=1	E	Automatic hyphenation is turned on for $N \geq 1$, or off for $N=0$. If $N=2$, *last* 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.
.hc c	\%	\%	E	Hyphenation indicator character is set to c or to the default \%. The indicator does not appear in the output.
.hw word1 ...		ignored	-	Specify hyphenation points in words with imbedded minus signs. Versions of a word with terminal s are

implied; i. e. *dig—it* implies *dig—its*. 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.

Request Form	Initial Value	If No Argument	Notes	Explanation
.tl *'left'center'right'*	-	-	The strings *left*, *center*, and *right* 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.
.pc *c*	%	off	-	The page number character is set to *c*, or removed. The page-number register remains %.
.lt ±*N*	6.5 in	previous	E,m	Length of title set to ±*N*. The line-length and the title-length are *independent*. 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).

Request Form	Initial Value	If No Argument	Notes	Explanation
.nm ±*N M S I*		off	E	Line number mode. If ±*N* is given, line numbering is turned on, and the next output line numbered is numbered ±*N*. Default values are *M*=1, *S*=1, and *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.
.nn *N*	-	*N*=1	E	The next *N* 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.

Request Form	Initial Value	If No Argument	Notes	Explanation
.if *c anything*	-	-		If condition *c* true, accept *anything* as input; in multi-line case use \\{*anything*\\}.
.if !*c anything*	-	-		If condition *c* false, accept *anything*.
.if *N anything*	-		u	If expression *N* > 0, accept *anything*.
.if !*N anything*	-		u	If expression *N* ⩽ 0, accept *anything*.
.if '*string1*'*string2*' *anything*	-		-	If *string1* identical to *string2*, accept *anything*.
.if !'*string1*'*string2*' *anything*	-		-	If *string1* not identical to *string2*, accept *anything*.
.ie *c anything*	-		u	If portion of if-else; all above forms (like if).
.el *anything*	-	-		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.

Request Form	Initial Value	If No Argument	Notes	Explanation
.ev *N*	*N*=0	previous	-	Environment switched to environment $0 \leqslant N \leqslant 2$. Switching is done in push-down fashion so that restoring a previous environment *must* 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*.

Request Form	Initial Value	If No Argument	Notes	Explanation
.rd *prompt*	-	*prompt*=BEL	-	Read insertion from the standard input until two newlines in a row are found. If the standard input is the user's keyboard, *prompt* (or a BEL) is written onto the user's terminal. **rd** behaves like a macro, and arguments may be placed after *prompt*.
.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

Request Form	Initial Value	If No Argument	Notes	Explanation
.so *filename*		-	-	Switch source file. The top input (file reading) level is switched to *filename*. 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 *filename*		end-of-file	-	Next file is *filename*. The current file is considered ended, and the input is immediately switched to *filename*.
.pi *program*		-	-	Pipe output to *program* (NROFF only). This request must occur *before* any printing occurs. No arguments are transmitted to *program*.

20. Miscellaneous

Request Form	Initial Value	If No Argument	Notes	Explanation
.mc *c N*	-	off	E,m	Specifies that a *margin* character *c* appear a distance N to the right of the right margin after each non-empty text line (those 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 em in TROFF. The margin character used with this paragraph was a 12-point box-rule.

.tm *string*	-	newline	-	After skipping initial blanks, *string* (rest of the line) is read in *copy mode* and written on the user's terminal.
.ig *yy*	-	*.yy*=..	-	Ignore input lines. **ig** behaves exactly like **de** (§7) except that the input is discarded. The input is read in *copy mode*, and any auto-incremented registers will be affected.
.pm *t*	-	all	-	Print macros. The names and sizes of all of the defined macros and strings are printed on the user's terminal; if *t* is given, only the total of the sizes is printed. The sizes is given in *blocks* 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
.ie \\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 ev1
.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 foot-note, **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 redivert 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 ¼ 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 fonts available from that company.

Times Roman

abcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZ
1234567890
! $ % & ( ) ' ' * + − . , / : ; = ? [ ] |
● □ — - _ ¼ ½ ¾ fi fl ff ffi ffl ° † ' ¢ ® ©

Times Italic

abcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZ
1234567890
*! $ % & ( ) ' ' * + − . , / : ; = ? [ ] |*
● □ — - _ ¼ ½ ¾ fi fl ff ffi ffl ° † ' ¢ ® ©

Times Bold

abcdefghijklmnopqrstuvwxyz
ABCDEFGHIJKLMNOPQRSTUVWXYZ
1234567890
! $ % & ( ) ' ' * + − . , / : ; = ? [ ] |
● □ — - _ ¼ ½ ¾ fi fl ff ffi ffl ° † ' ¢ ® ©

Special Mathematical Font

" ' \ ^ _ ` ˜ / < > { } # @ + − = *
$\alpha \beta \gamma \delta \epsilon \zeta \eta \theta \iota \kappa \lambda \mu \nu \xi o \pi \rho \sigma \varsigma \tau \upsilon \phi \chi \psi \omega$
$\Gamma \Delta \Theta \Lambda \Xi \Pi \Sigma \Upsilon \Phi \Psi \Omega$
$\sqrt{} \geqslant \leqslant \equiv \sim \simeq \neq \rightarrow \leftarrow \uparrow \downarrow \times \div \pm \cup \cap \subset \supset \subseteq \supseteq \infty \partial$
$\S \nabla \neg \int \propto \varnothing \in \ddagger$ ◼▶ ◀◼ ⓐ I O ⌈ | ⌉ ⌊ ⌋ ⌈ ⌉

Table II

Input Naming Conventions for ´, `, and −
and for Non-ASCII Special Characters

Non-ASCII characters and *minus* on the standard fonts.

Char	Input Name	Character Name	Char	Input Name	Character Name
´	´	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.

Char	Input Name	Character Name	Char	Input Name	Character Name
+	\(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	o	\(*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	A	\(*A	Alpha†
ι	\(*i	iota	B	\(*B	Beta†

Char	Input Name	Character Name
Γ	\(*G	Gamma
Δ	\(*D	Delta
E	\(*E	Epsilon†
Z	\(*Z	Zeta†
H	\(*Y	Eta†
Θ	\(*H	Theta
I	\(*I	Iota†
K	\(*K	Kappa†
Λ	\(*L	Lambda
M	\(*M	Mu†
N	\(*N	Nu†
Ξ	\(*C	Xi
O	\(*O	Omicron†
Π	\(*P	Pi†
P	\(*R	Rho†
Σ	\(*S	Sigma
T	\(*T	Tau†
Υ	\(*U	Upsilon
Φ	\(*F	Phi
X	\(*X	Chi†
Ψ	\(*Q	Psi
Ω	\(*W	Omega
√	\(sr	square root
	\(rn	root en extender
≥	\(>=	>=
≤	\(<=	<=
≡	\(==	identically equal
≈	\(~=	approx =
~	\(ap	approximates
≠	\(!=	not equal
→	\(->	right arrow
←	\(<-	left arrow
↑	\(ua	up arrow
↓	\(da	down arrow
×	\(mu	multiply
÷	\(di	divide
±	\(+-	plus-minus
∪	\(cu	cup (union)
∩	\(ca	cap (intersection)
⊂	\(sb	subset of
⊃	\(sp	superset of
⊆	\(ib	improper subset
⊇	\(ip	improper superset
∞	\(if	infinity
∂	\(pd	partial derivative
∇	\(gr	gradient
¬	\(no	not
∫	\(is	integral sign
∝	\(pt	proportional to
∅	\(es	empty set
∈	\(mo	member of

Char	Input Name	Character Name
|	\(br	box vertical rule
‡	\(dd	double dagger
☞	\(rh	right hand
☜	\(lh	left hand
☻	\(bs	Bell System logo
|	\(or	or
O	\(ci	circle
⌠	\(lt	left top of big curly bracket
⌡	\(lb	left bottom
⌐	\(rt	right top
⌐	\(rb	right bot
{	\(lk	left center of big curly bracket
}	\(rk	right center of big curly bracket
|	\(bv	bold vertical
⌊	\(lf	left floor (left bottom of big square bracket)
⌋	\(rf	right floor (right bottom)
⌈	\(lc	left ceiling (left top)
⌉	\(rc	right ceiling (right top)

Summary of Changes to N/TROFF Since the October 11, 1976 Manual

May 15, 1977

Options

– h (Nroff only) Output tabs used during horizontal spacing to speed output as well as reduce output byte count. Device tab settings 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 previously obtained from the ".j" register (see below).

.so name The contents of file "name" will be interpolated at the point the "so" is encountered. Previously, the interpolation was done upon return to the file-reading input level.

New Requests

.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.

.fz F N forces font "F" to be in size N. N may have the form N, + N, or – N. For example,
 .fz 3 – 2
will cause an implicit \s–2 every time font 3 is entered, and a corresponding \s+2 when it is left. Special font characters occurring during the reign of font F will have the same size modification. If special characters are to be treated differently,
 .fz S F N
may be used to specify the size treatment of special characters during font F. For example,
 .fz 3 – 3
 .fz S 3 – 0
will cause automatic reduction of font 3 by 3 points while the special characters would not be affected. Any "fp" request specifying a font on some position must precede "fz" requests relating to that position.

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. A number representing the current adjustment mode and type. Can be saved and later given to the "ad" request to restore a previous mode.

.P Read-only. 1 if the current page is being printed, and zero otherwise.

.L Read-only. Contains the current line-spacing parameter ("ls").

c. General register access to the input line-number in the current input file. Contains the same value as the read-only ".c" register.

I. NEW OPTIONS TO NROFF/TROFF

A. Compacted Macros

It is now possible to obtain a pre-processed version of a macro package, called "compacted macros." A compacted macro package is loaded (initialized) more efficiently, therefore faster, than its non-compacted version, but is functionally equivalent. However, it is not possible to perform a .so to a compacted macro package.

A compacted macro package consists of three files:

    cmp.[nt].d.*name*
    cmp.[nt].t.*name*
    ucmp.[nt].*name*

where n stands for an *nroff*- compacted package and the (first) t stands for a *troff*- compacted package.

Macro packages are compacted with the new −k option (see below). The macro source file must be arranged so that all compactable information is at the beginning of the file and all non-compactable material is at the end of the file. A .co request must separate the two sections of the source file. This means that a typical macro package will be arranged as below in order to be compacted:

    Compactable material
        .
        .
        .
    .co
    Non-compactable material
        .
        .
        .

The files d.*name* and t.*name* in the current directory will be used as output files for compacted information; cmp.[nt]. must be prepended to the names of these files, then the files cmp.[nt].[dt].*name* must be moved to the directory /usr/lib/macros. The file ucmp.[nt].*name* must be generated by hand. That is, the part of the source file containing the non-compactable material must be copied to /usr/lib/macros/ucmp.[nt].*name*.

Information that should not be compacted includes requests that interact with the command line (e.g., registers to be initialized from the command line) or refer to run-dependent information (e.g., mo, dy, and yr strings). If all information is compactable, the file /usr/lib/macros/ucmp.[nt].*name* may be empty, but it must exist.

NOTE: the *nroff* version of a macro package must be generated by *nroff*, and the *troff* version must be generated by *troff*. These files must reside in /usr/lib/macros in order to be referenced by *nroff* and *troff*.

There are two new options to *nroff* and *troff* for compacted macros:

    − c*name*   Use the compacted version of this macro package. For example:

            nroff − T300 − mm file

    and

            nroff − T300 − cm file

    are equivalent, with the exception that the first calls the non-compacted macros
    and the second calls the compacted version. When this option is used, an
    error message will be printed if a macro package is not available in compacted
    form.

-k*name* Compact this invocation of *nroff* or *troff*, where *name* is the name of the output file to be produced. The file to be compacted is given as an input file, as in the following example:

> nroff **-k***name* macrofile

The files t.*name* and d.*name* in the current directory will be used as output files; *nroff* and *troff* will build these files, they need not previously exist. The macro file must contain a .co request; only lines before this request will be compacted. Both **-k** and .co are necessary, if no .co is found in the file the **-k** is ignored, and if no **-k** appears on the command line the .co is ignored.

WARNING: the **-c***name* option will only succeed if *name* represents the name of a compacted macro package. This compacted package must also have been produced by the same version of *nroff* or *troff* that attempts to use the package. In other words, if *nroff* or *troff* is changed and recompiled, all compacted macro packages must be regenerated before they can be used.

B. TROFF Stop Message

A new command line option to *troff*, **-S** will cause the output of the message *pages finished* when *troff* is waiting for the user to restart the typesetter after pausing. (A pause is specified on the command line via the **-s** option.) The message *troff finished* will be output when the job has been completed.

NOTE: **-S** works only when **-s** is also used. For example, troff **-s5 -S ...** will cause *troff* to stop every five pages and print the message *pages finished*, with the message *troff finished* at the end of the job.

II. NEW NROFF/TROFF REQUEST

> .co Specify the point in the macro file at which compaction ends. When **-k***name* is called on the command line, all lines in file *name* before the .co request will be compacted.

III. FEATURES SCHEDULED TO BE REMOVED

The following features will be removed from future versions of *nroff/troff*:

> .li The transparent input mode request.
>
> **±N** *Nroff/troff* command line options: start(+) printing from page *N* or stop(-) printing at page *N*.

IV. NEW READ-ONLY REGISTERS

> .R The number of number registers that remain available for use.
>
> .b The emboldening factor of the current font.

V. NEW ESCAPE SEQUENCE

> \jx \j(xx mark the current horizontal output position in register *x* or *xx*

January 1980

ADDENDUM

The following changes to nroff/troff are effective with the
UNIX[TM] System III release.

Changes to nroff/troff fall into and are described under two
categories: enchancements and bug fixes.

NROFF/TROFF ENHANCEMENTS

Nroff Bold Font Simulation

By default, text which is in bold font in troff, is
emboldened (default is .db 3 3) in nroff. The user may
suppress this overstrike feature by using the nroff -u
command line option.

EXAMPLE:

                 nroff -u [options] file

However, -u does not suppress bold produced by terminal
hardware.

The -u option may take an argument which is the emboldening
factor for bold font.

EXAMPLE:

                 nroff -u10 [options] file

makes the emboldening factor 10, which is 10 overstrikes.

The degree of emboldenment may also be controlled by setting
the second argument of the .bd request.

EXAMPLE:

                 .bd 3 0

turns off emboldenment while

EXAMPLE:

.bd 3 3

overstrikes 3 times (default).

Terminals which have hardware bold capability, (such as the DASI 300s or the DTC 382) use that feature rather than emboldenment.

The old -u command line option has been replaced by the new -y command line option.

Escape and BEL Characters

Escape and BEL characters are valid characters in nroff and are accepted and passed to the output of the formatter. Troff continues to absorb these characters.

New Escape Sequence \gx

New escape sequences are available, \gx and \g(xx , which return the format of number register x(xx) in a form suitable for the .af request. If this is the first reference to the register, then the \g sequence returns the null string.

New Request for Command Execution (.!)

.! program [arg1] [arg2] [arg3] ...

A new nroff request, .! , has been added. This allows arbitrary command execution from within nroff/troff text, where the standard output of the command is placed at the point in the text where the command was executed. Standard input for the called program is closed.

Arguments passed to .! are passed as arguments to the program, therefore, quoting conventions are recognized.

Driving Tables

New driving tables have been added.

-TX	text train line printers
-T2631	HP2631A line printer
-T2631-e	HP2631A line printer, expanded
-T2631-c	HP2631A line printer, compressed

The -TX driving table includes special escape sequences denoting EBCDIC character codes that may be used by

postprocessors to nroff.

All driving tables have been modified to include escape
sequences for italic on and off.

The 450-12-8 driving table has been removed.

Hyphenation

Hyphenation is off by default.

Page Stop Option

The -S command line option is removed.

The -s (page stop) command line option now works correctly
and causes diagnostics to be output between pages. A bell
(ASCII 007) is output for nroff and the message "page stop"
is output for troff.

Transparent Line Mode

The undocumented and unsupported .li (transparent line mode)
request has been removed.

Environment Blocks

The environment block no longer depends on the order of
definition to be the order of real storage for environment
variables. The environment block is now a structure, and is
a portable feature. (Note: this feature is invisible to
users.)·

Fonts, Terms, and Suftab

For portability purposes, fonts, terms, and suftab are now
pure data files, not a.out/assembler formats.

Invocation of Multiple Macro Packages

Nroff/Troff accepts any number of different macro packages
per run.

EXAMPLE:
                      nroff -mm -mosd file

Each macro package, however, should only be invoked once.

New Troff Option for Output Device (-Tname)

       troff -T[name] [options] file

The troff -T option specifies the output device. By
default, the output device is the C/A/T. For UNIX Release
3.0, this is the only supported troff output device. (It is
expected that future releases may support other devices via
this mechanism.)

New Compacted Macros Option (-cname)

Compacted macros may be called by replacing the -m option
with -c.

EXAMPLE:
                        nroff -cm -cosd file

If compacted macros do not exist for a particular package,
the non-compacted macro package is called automatically.
Therefore, -c should always be used.

Print Only Page Option

The +n/-n option for nroff/troff to print specific pages is
removed.

Use the -o option documented in the "Nroff/Troff User's
Manual".

NROFF/TROFF BUG FIXES

Failed if Statements

A failed .if request, including multi-line cases, ignores
the anything part, where anything represents what is
conditionally accepted. For example, automatic incrementing
or decrementing of registers does not occur, and strings are
not expanded.

Control Characters

A control character may always be separated from a request
or macro name by white space (space and/or tabs). This
includes end brackets to .de and .ig.

.so Request

Nroff no longer loses the end of the input block when a .so
is done with nroff's buffer pointers pointing at the end of
the input buffer.

This particular problem formerly caused loss of input text
after the .so.

Continuous Underlining of Special Characters in Diversions

Continuous underlining (.cu) of special characters within a
diversion now works properly and the word buffer size has
been increased to allow approximately three lines of
continuous underlining.

Compacted Macros

Compacted macros have been redesigned for reliability and
simplicity. The only formatter characteristics that may be
compacted are: number registers, strings, diversions,
macros, traps, and the translation table (.tr).

Compacted macros are used by default in the following macro
package shell commands: mm, mmt, and man.

Thetas Fixed for -T37 Driving Table

The upper and lower case thetas in the -T37 driving table
have been fixed (they were reversed).

Print Only Page Option with Compacted Macros

The -o command line option works properly with compacted
macros.
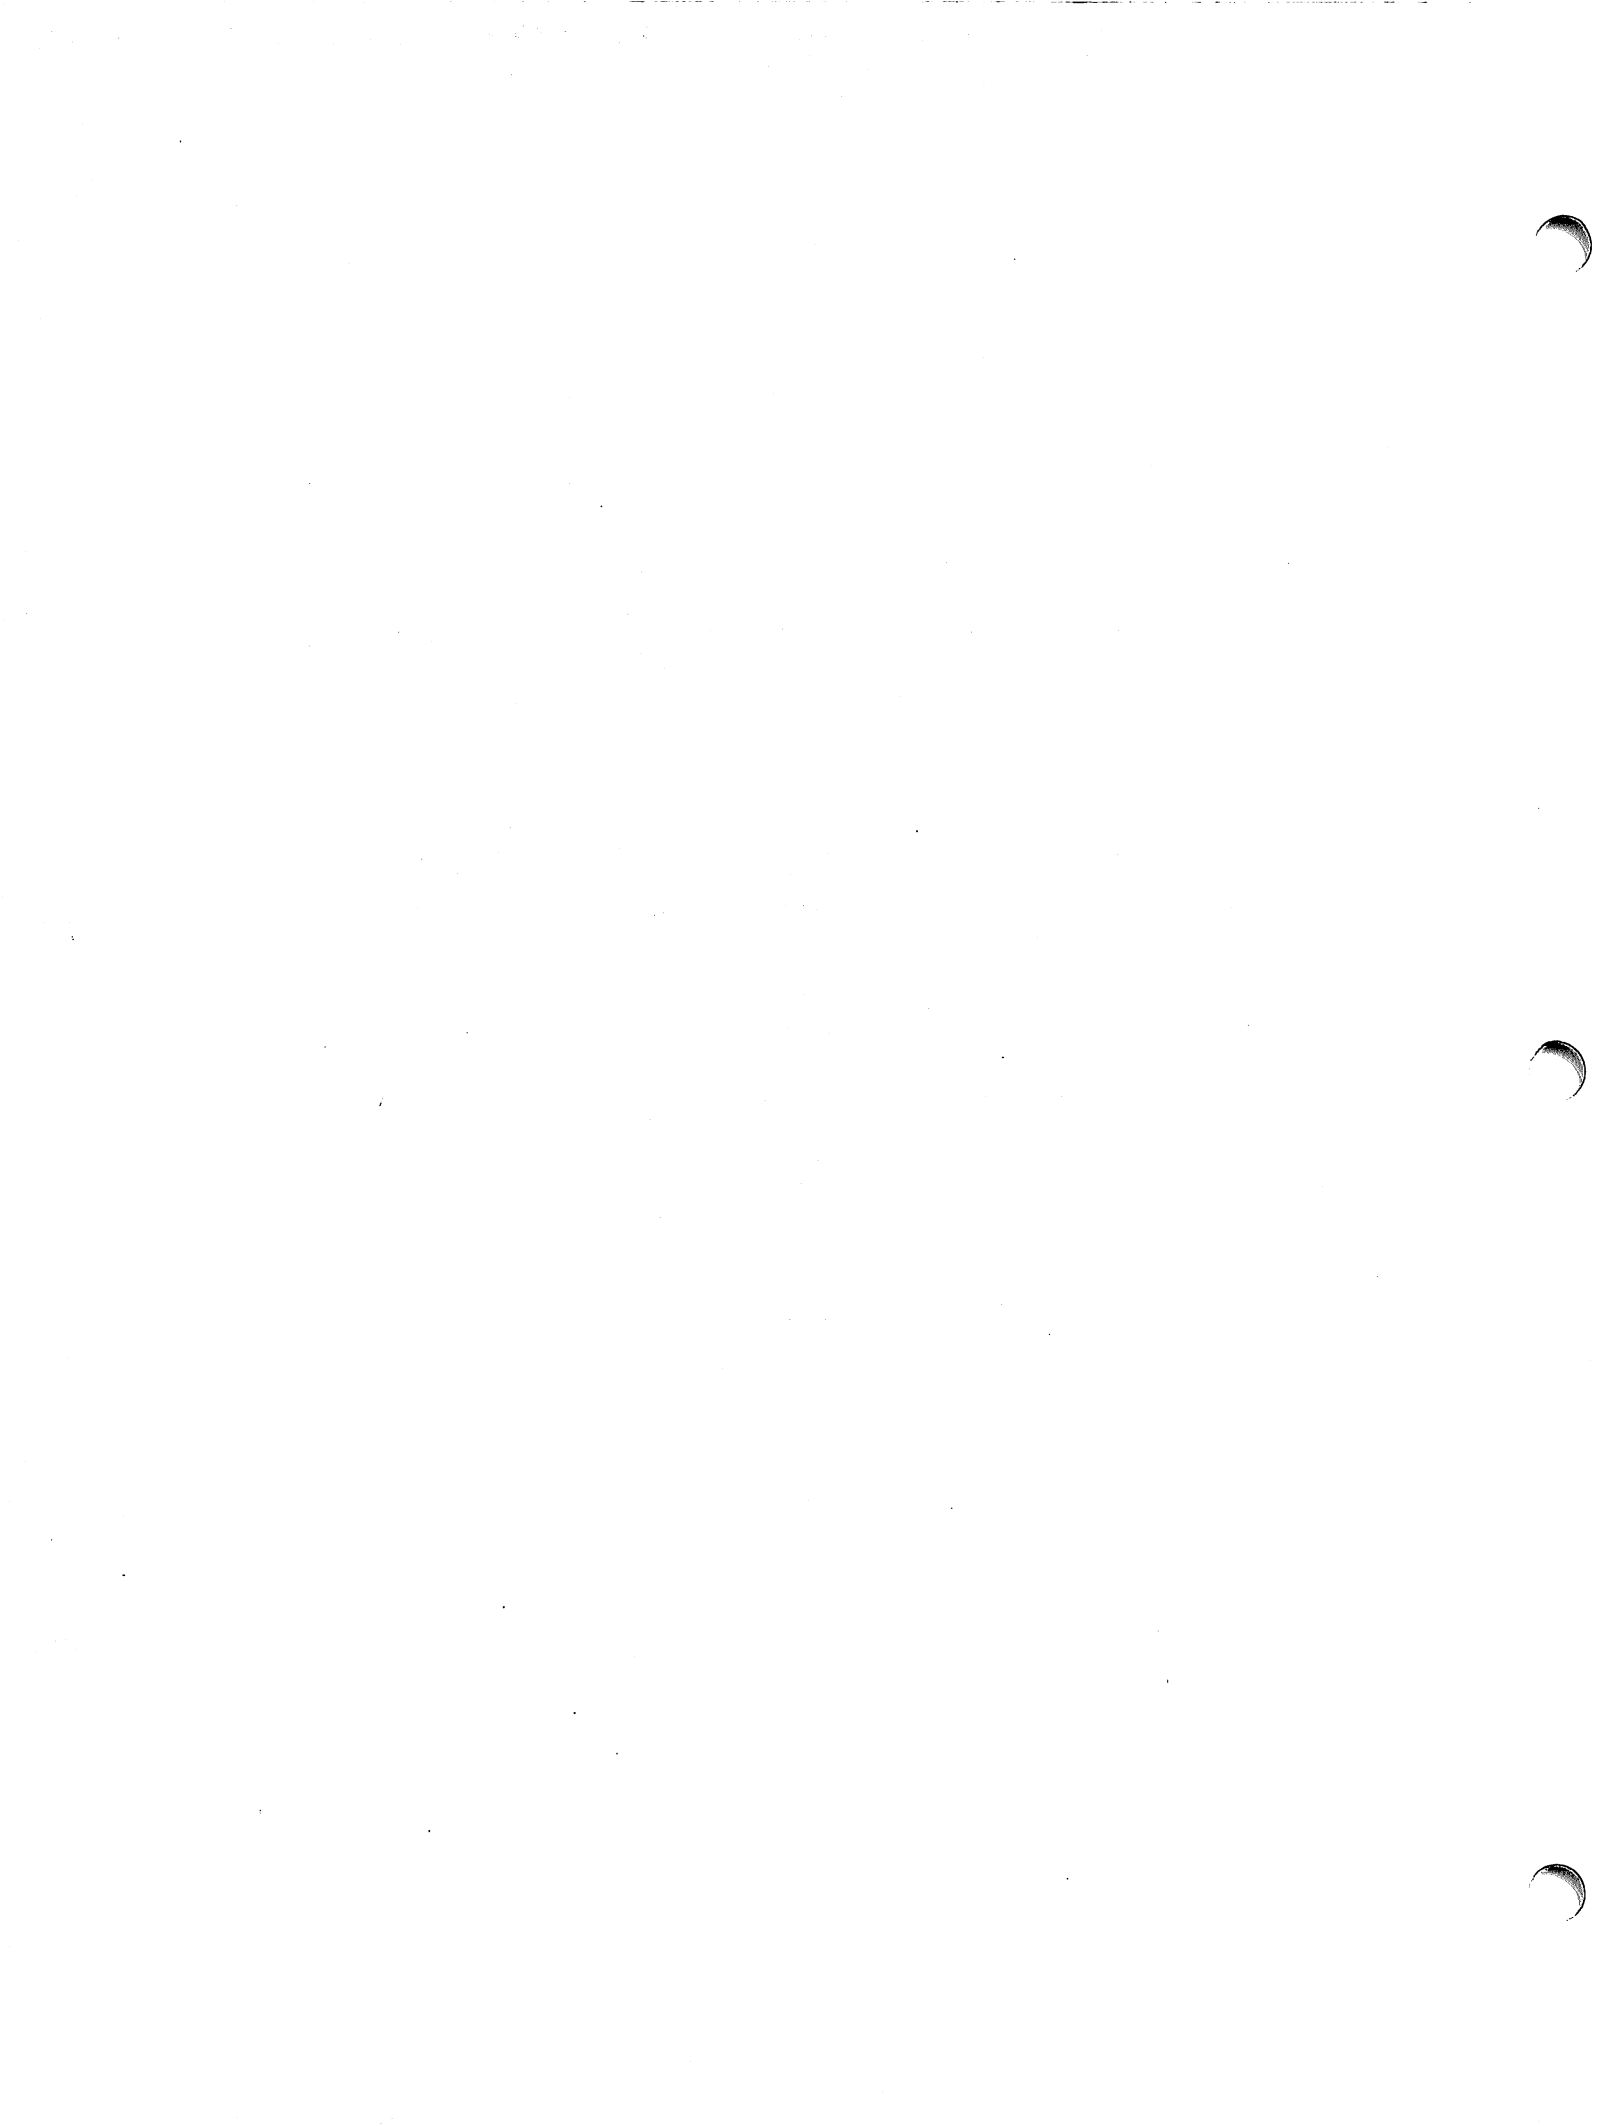
A TROFF Tutorial

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

troff is a text-formatting program for driving the Graphic Systems photo-typesetter 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. In most respects, the UNIX formatter **nroff** is identical to **troff**, so this document also serves as a tutorial on **nroff**.

August 4, 1978

A TROFF Tutorial

Brian W. Kernighan

Bell Laboratories
Murray Hill, New Jersey 07974

1. 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

$$Area = \pi r^2$$

you have to type

Area = \(*p\fIr\fR\|\s8\u2\d\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 ½ 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.

abcdefghijklmnopqrstuvwxyz 0123456789
ABCDEFGHIJKLMNOPQRSTUVWXYZ
abcdefghijklmnopqrstuvwxyz 0123456789
ABCDEFGHIJKLMNOPQRSTUVWXYZ
abcdefghijklmnopqrstuvwxyz 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 .ft3 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:

\(*S	Σ
(	(
\(*a	α
\(mu	$\times$
\(*b	β
)	)
\(−>	$\rightarrow$
\(if	∞

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; adveniat regnum tuum; fiat voluntas 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 $= \pi r^2$' 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

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 ln)
.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 \v, like \u and \d, should always balance upward vertical motion in a line with the same amount in the downward direction.

Arbitrary horizontal motions are also available — \h is quite analogous to \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 \h is used with the 'width function' \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 \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 \(blank), which is an unpaddable character the width of a space, \|, which is half that width, \^, which is one quarter of the width of a space, and \&, which has zero width. (This last one is useful, for example, in entering a text line which would otherwise begin with a '.'.)

The command \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

système téléphonique

The accents are \(ga and \(aa, or \' and \'; remember that each is just one character to **troff**.

You can make your own overstrikes with another special convention, \z, the zero-motion command. \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 \o, it centers the characters on the widest, and there can be no horizontal or vertical motions, so \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 semi-colon that looks like

; instead of ; or ;

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:

{[x]}

by typing in only this:

.sp
\b'\(lt\(lk\(lb' \b'\(lc\(lf' x \b'\(rc\(rf' \b'\(rt\(rk\(rb'

troff also provides a convenient facility for drawing horizontal and vertical lines of arbitrary length with arbitrary characters. \l'li' 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
    .ti +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'
.fo '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 −1i 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 nth 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 .$.

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\fl\\$1\h'-\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 1          \" 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 Computing Science Technical Report 54, 1976.

[2]   B. W. Kernighan, *A System for Typesetting Mathematics — User's Guide (Second Edition)*, Bell Laboratories Computing Science Technical Report 17, 1977.

[3]   M. E. Lesk, *TBL — A Program to Format Tables*, Bell Laboratories Computing Science Technical Report 49, 1976.

[4]   M. E. Lesk, *Typing Documents on UNIX*, Bell Laboratories, 1978.

[5]   J. R. Mashey and D. W. Smith, *PWB/MM — 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.

ff	\\(ff	fi	\\(fi	fl	\\(fl	ffi	\\(Fi	ffl	\\(Fl
_	\\(ru	—	\\(em	¼	\\(14	½	\\(12	¾	\\(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	ş	\\(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 \\ # @

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 o π ρ σ τ υ φ χ ψ ω
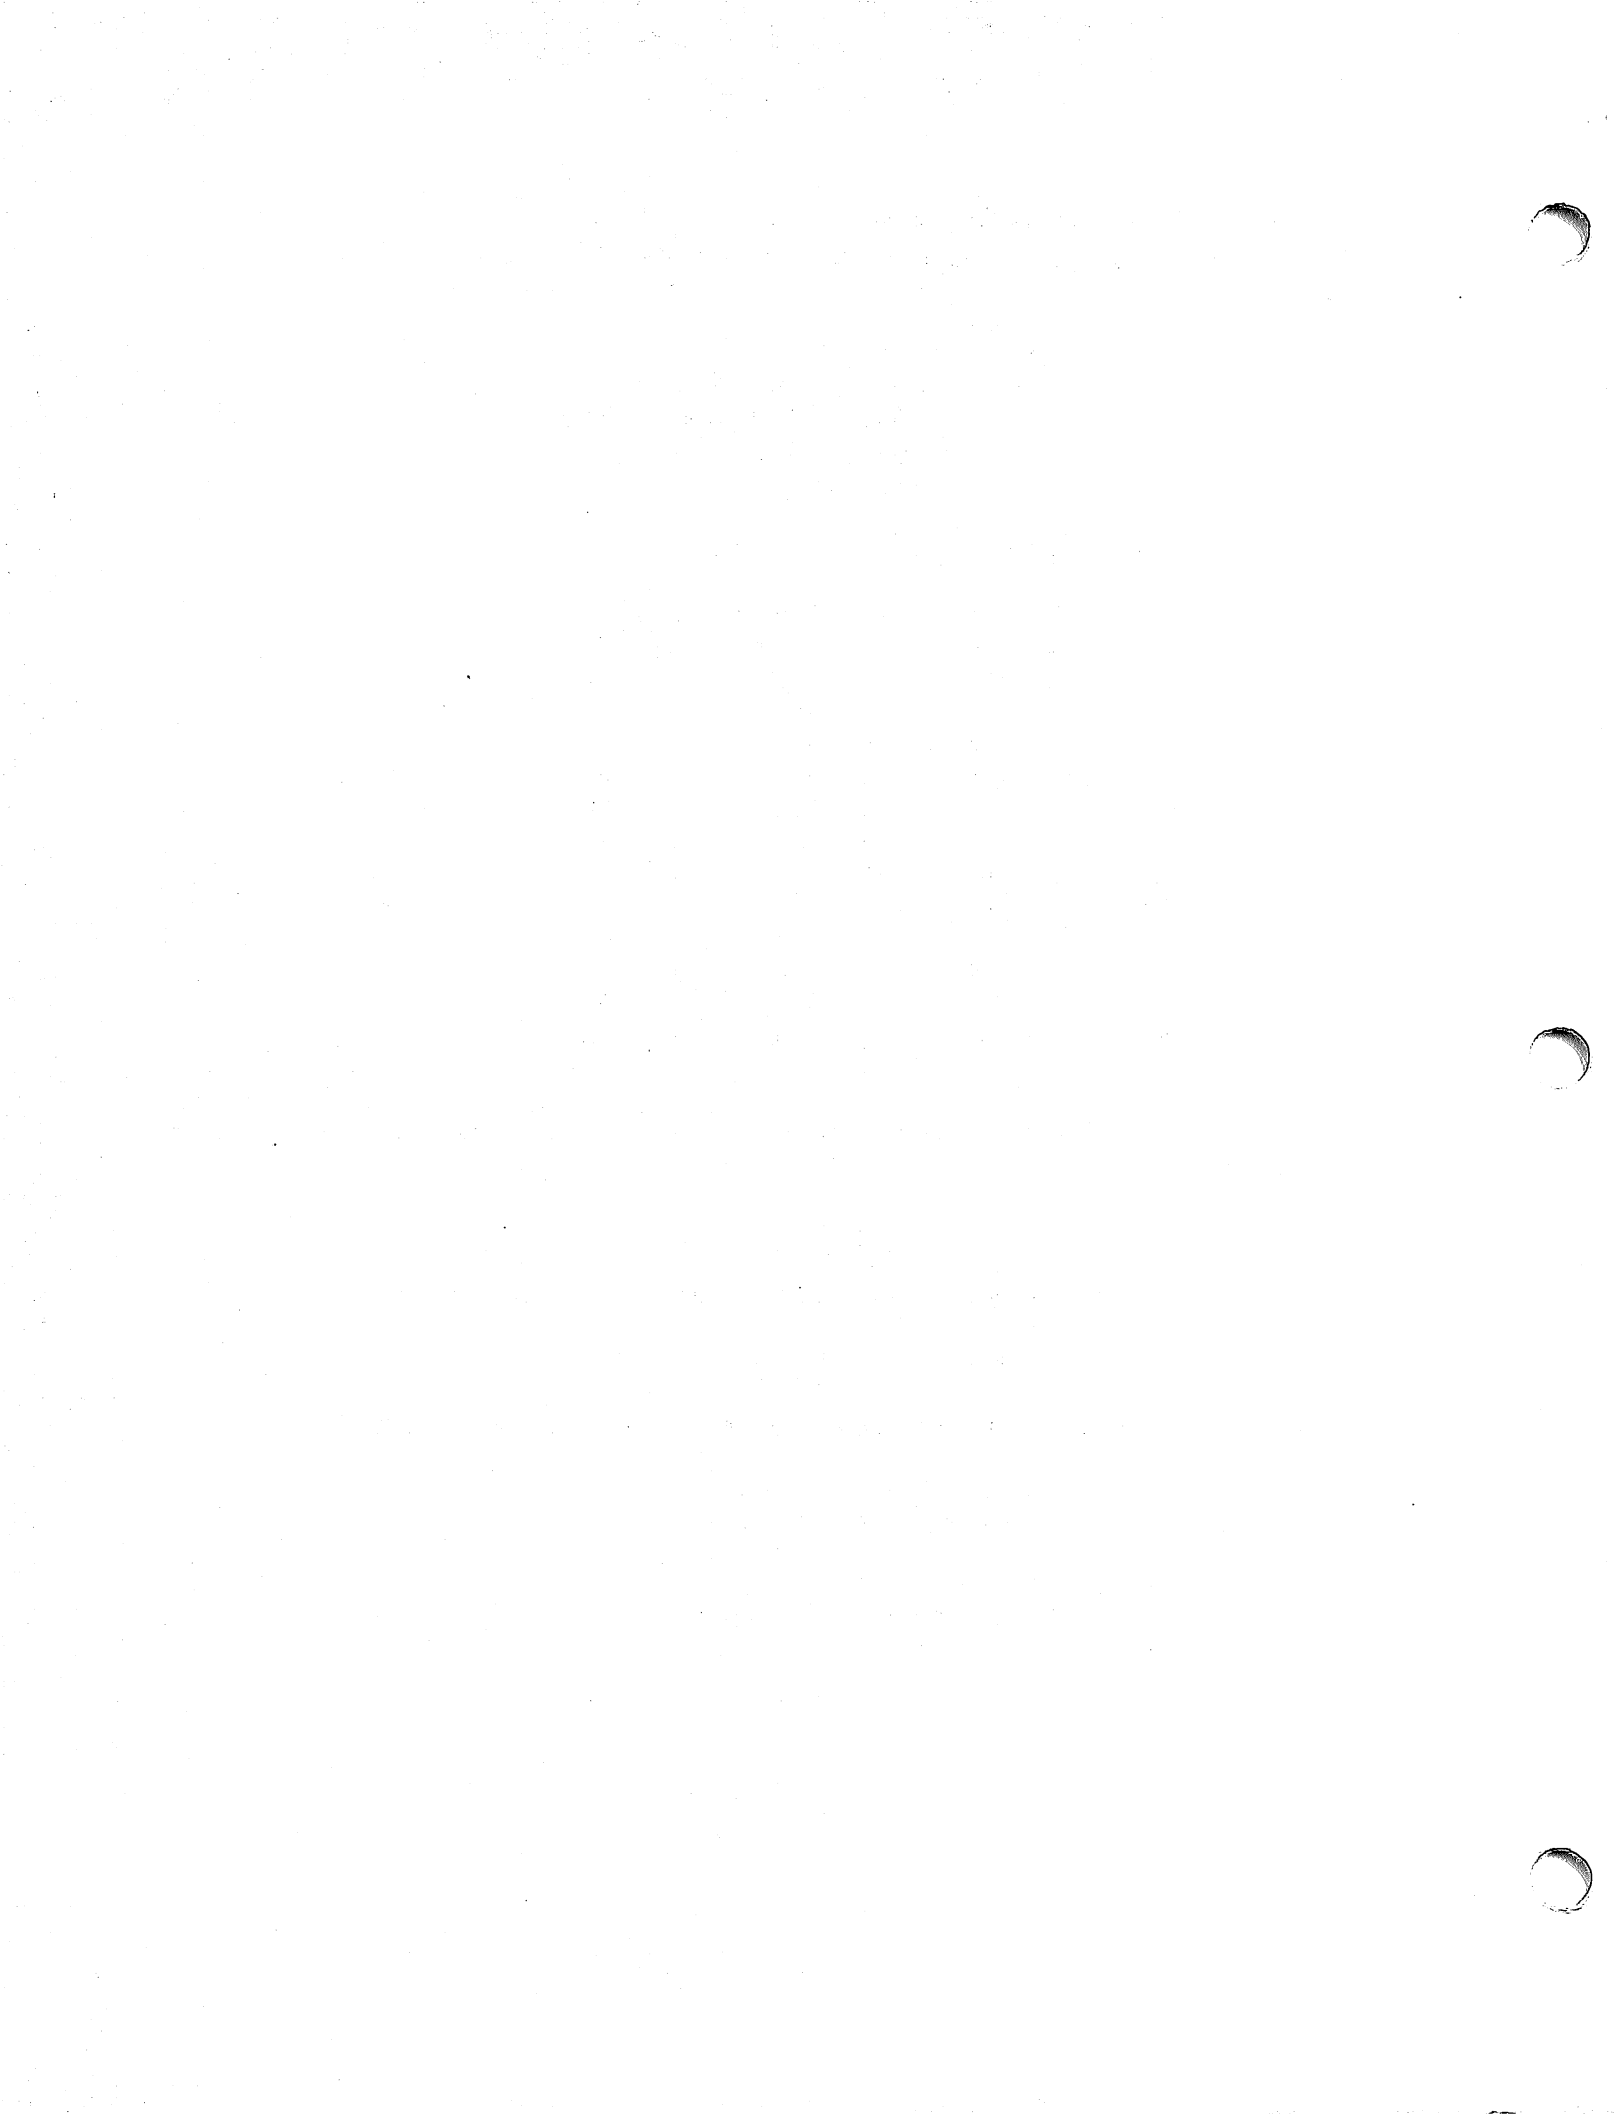
A B G D E Z Y H I K L M N C O P R S T U F X Q W
A B Γ Δ E Z H Θ I K Λ M N Ξ O Π P Σ T Y Φ X Ψ Ω

PWB/MM

Programmer's Workbench

Memorandum Macros

D. W. Smith
J. R. Mashey
E. C. Pariser (January 1980 Reissue)

PWB/MM
Programmer's Workbench Memorandum Macros

LIST OF FIGURES

PWB/MM—Programmer's Workbench Memorandum Macros

D. W. Smith
J. R. Mashey
E. C. Pariser (January 1980 Reissue)

Bell Laboratories
Piscataway, New Jersey 08854

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, 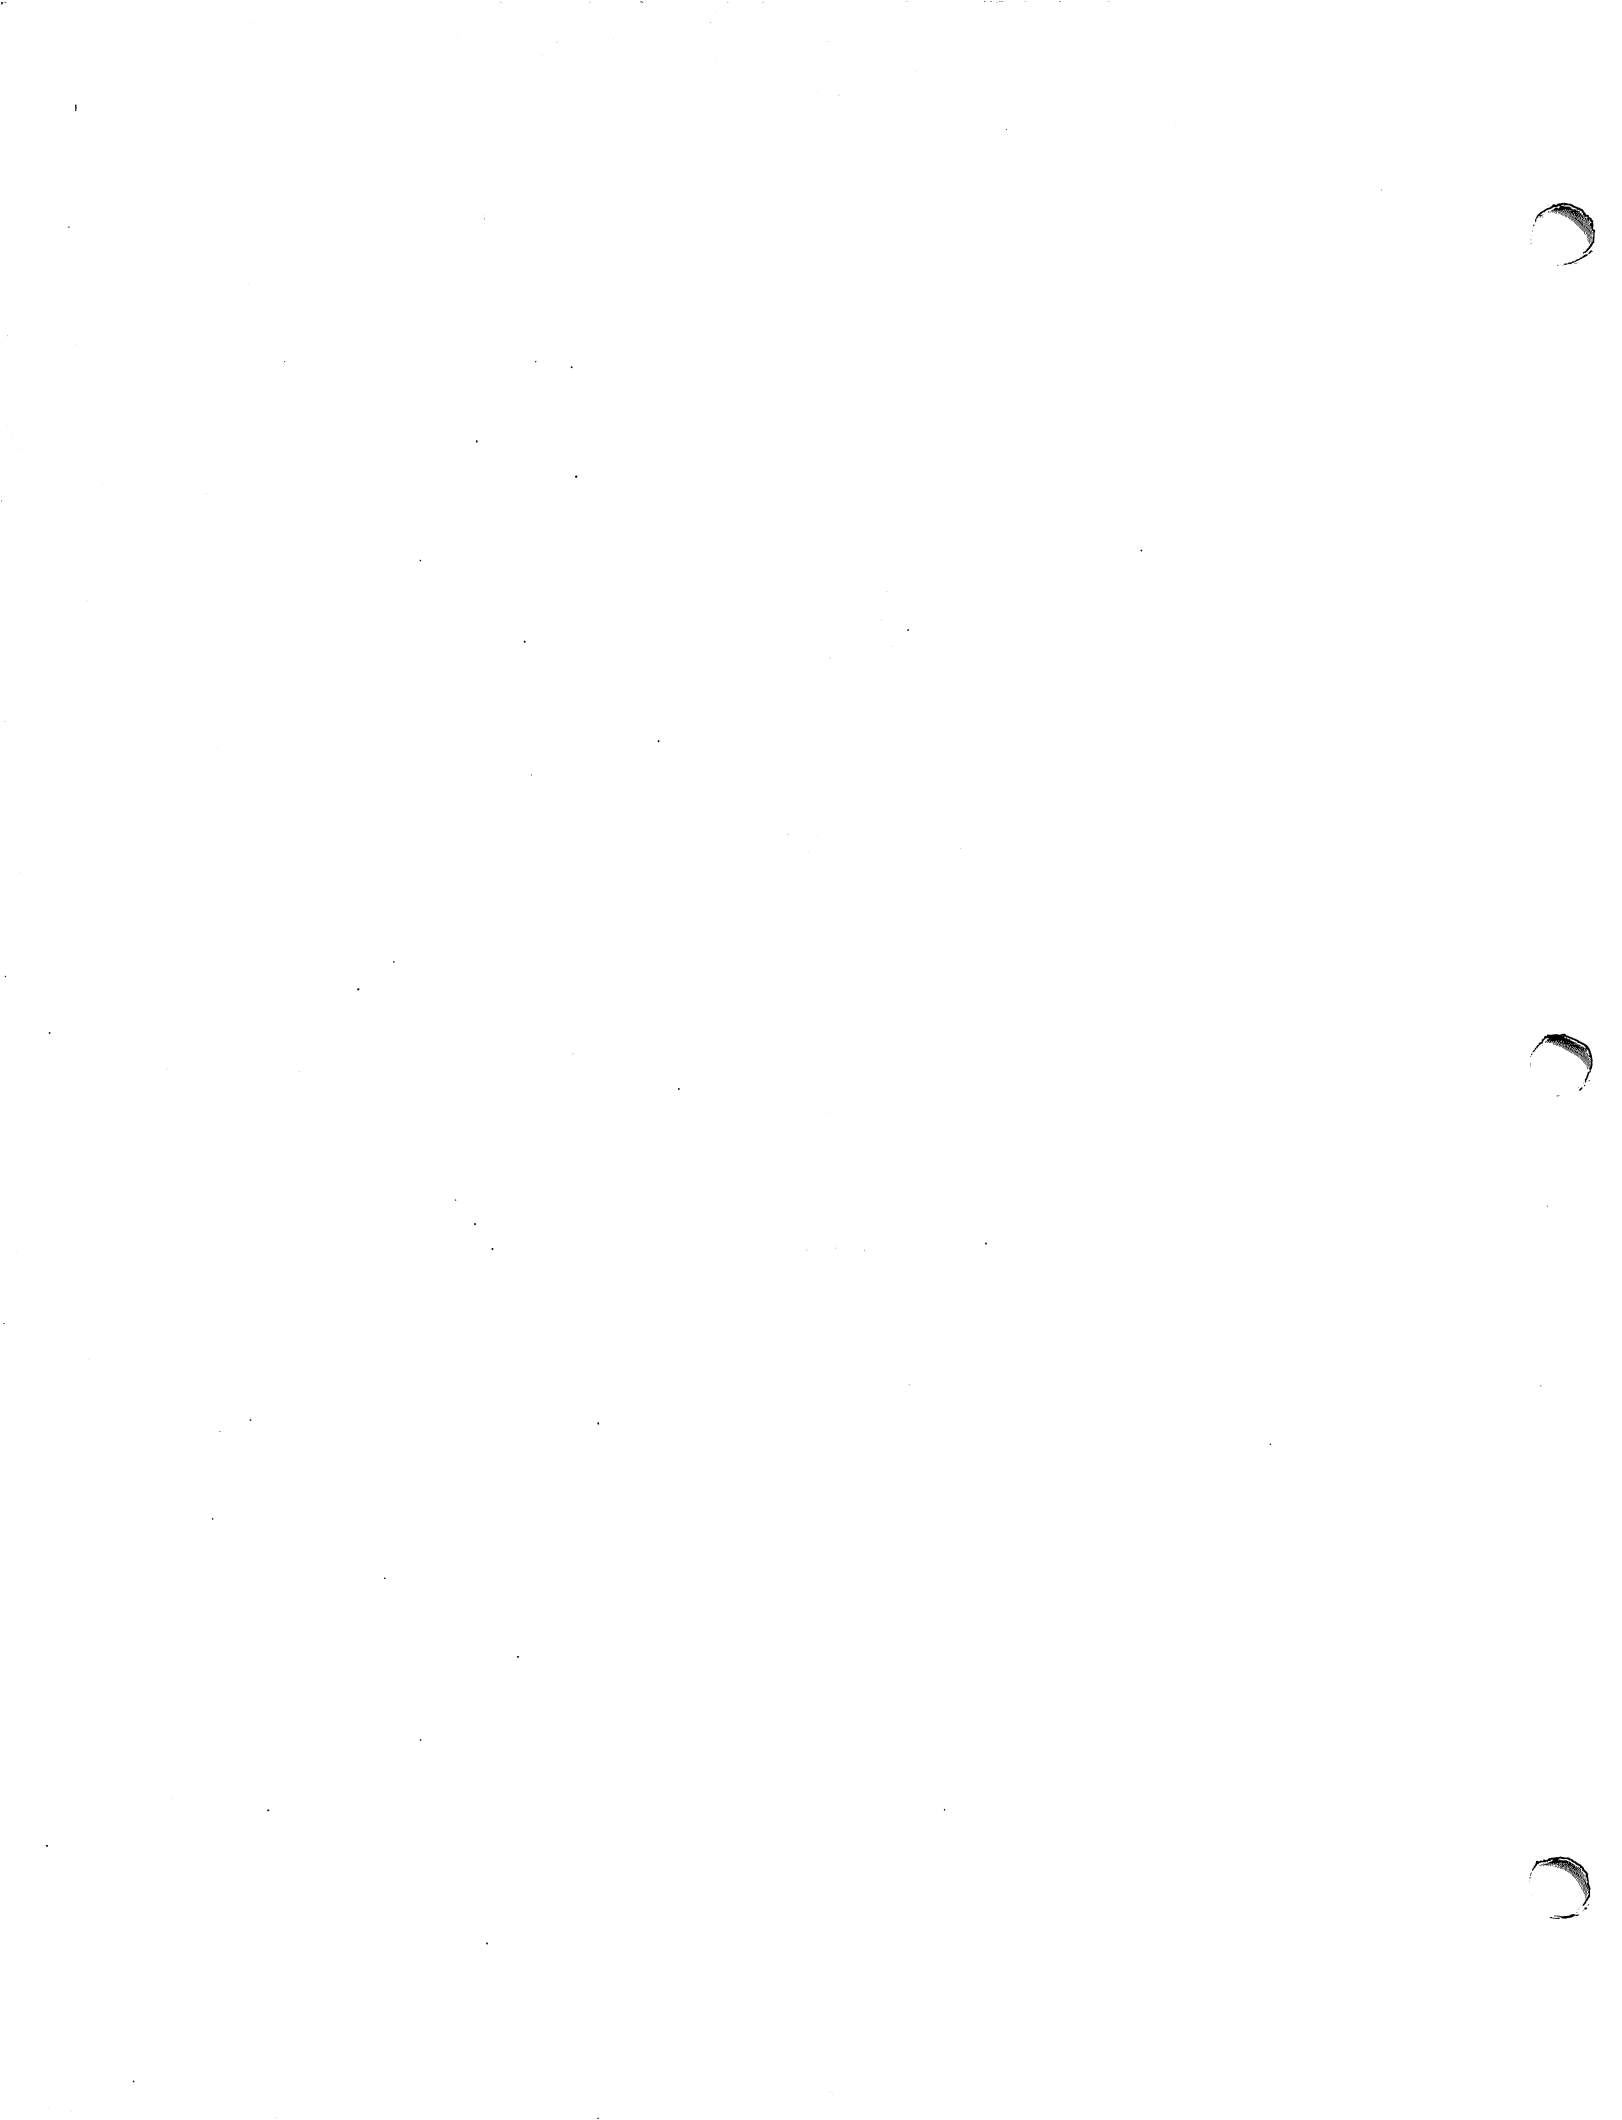consistent, and flexible tool for producing many common types of documents. Although PWB/UNIX provides other macro packages for various *specialized* formats, PWB/MM has become the standard, general-purpose macro package for most documents.

PWB/MM can be used to produce:

- Letters.
- Reports.
- Technical Memoranda.
- Released Papers.
- 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].

---

† UNIX is a Trademark of Bell Laboratories.

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 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, e.g., title, author's name, date.

- *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. Five 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 document (10).

The existence and size of these four segments varies widely among different document types. Although a specific item (such as date, title, 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 January 22, 1980.

The current date can be replaced, e.g.:

.ds DT 01/01/79

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.
-E	the "-e" option of *nroff* is to be invoked.
-12	need 12-pitch mode. Be sure that the pitch switch on the terminal is set to 12.
-T300	output is to a DASI300 terminal. This is the *default* terminal type (unless STERM is set).
-T300-12	output is to a DASI300 in 12-pitch mode.
-T300s	output is to a DASI300S.
-T300S	output is to a DASI300S.
-T300s-12	output is to a DASI300S in 12-pitch mode.
-T300S-12	output is to a DASI300S in 12-pitch mode.
-T4014	output is to a Tektronix 4014.
-Thp	output is to a HP264x.

```
-T450    output is to a DASI450.
-T450-12 output is to a DASI450 in 12-pitch mode.
-Ttn     output is to a GE TermiNet 300.
-Ttn300  output is to a GE TermiNet 300.
-Tti     output is to a Texas Instrument 700 series terminal.
-T37   . output is to a TELETYPE® Model 37.
-T43     output is to a TELETYPE® Model 43.
```

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/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 *300*(1), *450*(1), *hp*(1), *col*(1), and *terminals*(7) 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 -T450 filename ...
    or nroff -T450 -h -mm filename ...

- DASI450 in 12-pitch, 6 lines/inch mode:

        mm -T450-12 filename ...
    or nroff -T450-12 -rW80 -rO3 -h -mm filename ...
    or nroff -T450-12 -rT1 -h -mm filename ...

- Hewlett-Packard HP264x CRT family:

        mm -Thp filename ...
    or nroff -h -mm filename ... | hp

- Any terminal incapable of reverse paper motion (GE TermiNet, Texas Instruments 700 series, etc.):

        mm -Ttn filename ...
    or nroff -mm filename ... | col

- Versatec printer (see *vp*(1) for additional details):

        vp [vp-options] "mm -rT2 -c filename ..."
    or vp [vp-options] "nroff -rT2 -mm filename ... | col"

Of course, *tbl*(1) and *eqn*(1)/*neqn*(1), 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*, either the -c option must be specified to *mm*(1), or the *nroff* output must be postprocessed by *col*(1). 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*(1) must be processed by the appropriate terminal filter (e.g., *300*(1)); *mm*(1) 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. When *n* is
         4, the *default* page header is suppressed; however a user-specified header is not affected.

n	*Page 1*	*Pages 2 ff.*
0	header	header
1	header replaces footer	header
2	no header	header
3	"section-page" as *footer*	
4	no header	no header unless PH defined

         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.

-rS*n*   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
         troff 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*.[1] 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 is 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/tmac.m
         remainder of text

---

1. For *nroff, k* is an *unscaled* number representing lines or character positions; for *troff, k* must be *scaled.*

In this case, one must *not* use the -mm flag (nor the *mm*(1) 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/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:

— 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 (").[2] Otherwise, it will be treated as several separate arguments.

---

2. A double quote (") is a *single* character that must not be confused with two apostrophes or acute accents (''), or with two grave accents (``).

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 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]

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* values[3] 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*(1) (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 *(EM for each text dash for both *nroff* and *troff*. This string generates an em dash (—) in *troff* and generates "--" in *nroff*. Note that the dash list (.DL) macros (5.3.3.3) automatically generate the em dashes for the list items.

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 Trademark String

A trademark string *(Tm is now available with PWB/MM. This places the letters "TM" one-half line above the text that it follows.

---

3. Every eight characters in *nroff*; every ½ inch in *troff*.

For example:

   The PWB/UNIX*(Tm User's Manual is available from the library.

   yields:
   The PWB/UNIX™ User's Manual is available from the library.

3.10 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 0, which always provides left-justified paragraphs. All paragraphs can be forced to be indented by inserting the following at the beginning of the document:

   .nr Pt 1

All paragraphs will be indented except after headings, lists, and displays if the following:

   .nr Pt 2

is inserted 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.

The number register *Ps* controls the amount of spacing between paragraphs. By default, *Ps* is set to 1, yielding one blank space (½ a vertical space).

☛ *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 .P is ignored. {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 $3 \leqslant n \leqslant 7$, 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 *Pt* {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.

For example, to cause a blank line (½ 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 *Pt*), 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 and is short enough to fit on a single line. If a heading is to be underlined, but 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}.

4.2.2.4.3 Heading Point Sizes.. The user may also specify the desired point size for each heading level
with the *HP* string (for use with *troff* only).

```
.ds HP [ps1] [ps2] [ps3] [ps4] [ps5] [ps6] [ps7]
```

By default, the text of headings (.H and .HU) is printed in the same point size as the body *except* that bold stand-alone headings are printed in a size one point smaller than the body. The string *HP*, similar to the string *HF*, can be specified to contain up to seven values, corresponding to the seven levels of headings. For example:

.ds HP 12 12 11 10 10 10 10

specifies that the first and second level headings are to be printed in 12-point type, with the remainder printed in 10-point. Note that the specified values may also be *relative* point-size changes, e.g.:

.ds HP +2 +2 -1 -1

If absolute point sizes are specified, then those sizes will be used regardless of the point size of the body of the document. If relative point sizes are specified, then the point sizes for the headings will be relative to the point size of the body, even if the latter is changed.

Omitted or zero values imply that the *default* point size will be used for the corresponding heading level.

☞ *Only the point size of the headings is affected. Specifying a large point size without providing increased vertical spacing (via .HX and/or .HZ) may cause overprinting.*

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.

Value	Interpretation
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 *Cl* 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 *Cl* (contents level) is saved and later displayed in the table of contents. The default value for *Cl* 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 o

☞ *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 (4.3), *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 *rlevel* is non-zero, this string contains the "heading mark." Two unpaddable spaces (to separate the *mark* from the *heading*) have been appended to this string. If *rlevel* 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 unpaddable spaces that will be used to separate the (run-in) *heading* from the following *text.* 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 *dlevel* 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 unpaddable 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 *H1* 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 H1 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 five 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 unpaddable 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 + xxx
    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 (paddable) 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 .LL.

.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 *Li* (initially 5)[4] 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.

Spacing at the beginning of the list and between the items can be suppressed by setting the *Ls* (list space) register. *Ls* is set to the innermost list level for which spacing *is* done. For example:

.nr Ls   0

specifies that no spacing will occur around *any* list items. The default value for *Ls* is 6 (which is the *maximum* list nesting level).

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}.[5] 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 *Li* for this list only. If *text-indent* is null, then the value of *Li* will be used.

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}.[6]

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.

---

4. Values that specify indentation must be *unscaled* and are treated as "character positions," i.e., as the number of *ens*.

5. Note that the "0001" format is *not* permitted.

6. So that, in the default case, the text of bullet and dash lists lines up with the first line of indented paragraphs.

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:

```
.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).[7] 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:

Type	Mark	Result
0	omitted	hanging indent
0	*string*	*string* is the mark
>0	omitted	arabic numbering
>0	one of:	automatic numbering or
	1, A, a, I, i	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:*

Type	Appearance
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.

---

7. The *mark-indent* argument is typically 0.

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).

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).[8] 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, 67890" 987654321
    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 5432 1Z-234

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

---

8. 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.

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 organizations have their own numbering schemes for memoranda, engineer's notes, etc. These numbers are printed 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 3322.11AB

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 7654321 77777777

This macro call is ignored in the released-paper and external-letter styles {6.6}.

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.[9]

Note that headings {4.2, 4.3} and displays {7} 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]

---

9. Values that specify indentation must be *unscaled* and are treated as "character positions," i.e., as the number of *ens*.

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:

Code	Value
.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 the string, 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.

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 codes[10] 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:

```
.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
```

---

10. The complete list is: AK, CP, CH, CB, DR, HO, IN, IH, MV, MH, PY, RR, RD, WV, and WH.

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 notations[11] 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 data[12] 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-9876/5432-JJJ/SPL-cen
```

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:

---

11. See [2], pp. 1.12–16

12. 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

Code	Notations
.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 "*string*"	Copy (*string*) 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 -rL*n* 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:[13] 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 or letter with the following meanings:

Code	Meaning
""	no indent
0 or L	no indent
1 or I	indent by standard amount
2 or C	center each line
3 or CB	center as a block

The *fill* argument is also an integer or letter and can have the following meanings:

Code	Meaning
""	no-fill mode
0 or N	no-fill mode
1 or F	fill mode

Omitted arguments are taken to be zero.

The standard amount of indentation is taken from the register Si, 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 Pi register {4.1}. Even though their initial values are the same, these two registers are independent of one another.

The display format value 3 (CB) centers the *entire display* as a block (as opposed to .DS 2 and .DF 2, which center each line *individually*). That is, all the collected lines are left-justified, and then the display is centered, based on the width of the longest line. This format *must* be used in order for the *eqn* (1)/*neqn* (1) "mark" and "lineup" feature to work with centered equations (see section 7.4 below).

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 Ds to 0.

---

13. Displays are processed in an environment that is different from that of the body of the text (see the .ev request in {9}).

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 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.

The user may exercise great control over the output positioning of floating displays through the use of two number registers, *De* and *Df*. When a floating display is encountered by *nroff* or *troff*, it is processed and placed onto a queue of displays waiting to be output. Displays are always removed from the queue and printed in the order that they were entered on the queue, which is the order that they appeared in the input file. If a new floating display is encountered and the queue of displays is empty, then the new display is a candidate for immediate output on the current page. Immediate output is governed by the size of the display and the setting of the *Df* register (see below). The *De* register (see below) controls whether or not text will appear on the current page after a floating display has been produced.

As long as the queue contains one or more displays, new displays will be automatically entered there, rather than being output. When a new page is started (or the top of the second column when in two-column mode) the next display from the queue becomes a candidate for output if the *Df* register has specified "top-of-page" output. When a display is output it is also removed from the queue.

When the end of a section (when using section-page numbering) or the end of a document is reached, all displays are automatically removed from the queue and output. This will occur before a .CS or .TC is processed.

A display is said to "fit on the current page" if there is enough room to contain the entire display on the page, or if the display is longer than one page in length and less than half of the current page has been used. Also note that a wide (full page width) display will never fit in the second column of a two-column document.

The registers, their settings, and their effects are as follows:

Values for *De* Register	
Value	Action
0	DEFAULT: No special action occurs.
1	A page eject will *always* follow the output of each floating display, so only one floating display will appear on a page and no text will follow it.
	NOTE: For any other values the action performed is for the value 1.

Values for *Df* Register	
Value	Action
0	Floating displays will not be output until end of section (when section-page numbering) or end of document.

Values for *Df* Register	
Value	Action
1	Output the new floating display on the current page if there is room, otherwise hold it until the end of the section or document.
2	Output exactly one floating display from the queue at the top of a new page or column (when in two-column mode).
3	Output one floating display on current page if there is room. Output exactly one floating display at the top of a new page or column.
4	Output as many displays as will fit (at least one), starting at the top of a new page or column. Note that if register *De* is set to 1, each display will be followed by a page eject, c:using a new top of page to be reached where at least one more display will be output. (This also applies to value 5, below.)
5	DEFAULT: Output a new floating display on the current page if there is room. Output at least one, but as many displays as will fit starting at the top of a new page or column.
	NOTE: For any value greater than 5 the action performed is for the value 5.

7.3 Tables

```
.TS [H]
global options;
column descriptors.
title lines
[.TH [N]]
data within the table.
.TE
```

The .TS (table start) and .TE (table end) macros make possible the use of the *tbl*(1) processor [11]. They are used to delimit the text to be examined by *tbl*(1) as well as to set proper spacing around the table. The display function and the *tbl*(1) delimiting function are independent of one another, however, so in order to permit one to keep together blocks that contain any mixture of tables, equations, filled and unfilled text, and caption lines the .TS-.TE block should be enclosed within a display (.DS-.DE). Floating tables may be enclosed inside floating displays (.DF-.DE).

The macros .TS and .TE also permit the processing of tables that extend over several pages. If a table heading is needed for each page of a multi-page table, specify the argument "H" to the .TS macro as above. Following the options and format information, the table heading is typed on as many lines as required and followed by the .TH macro. The .TH macro *must* occur. when ".TS H" is used. Note. that this is *not* a feature of *tbl*(1), but of the macro definitions provided by PWB/MM.

The table header macro .TH may take as an argument the letter N. This argument causes the table header to be printed only if it is the first table header on the page. This option is used when it is necessary to build long tables from smaller .TS H/.TE segments. For example:

```
.TS H
global options;
column descriptors.
Title lines
.TH
data
.TE
.TS H
global options;
column descriptors.
 Title lines
.TH N
data
.TE
```

will cause the table heading to appear at the top of the first table segment, and no heading to appear at the top of the second segment when both appear on the same page. However, the heading will still appear at the top of each page that the table continues onto. This feature is used when a single table must be broken into segments because of table complexity (for example, too many blocks of filled text). If each segment had its own .TS H/.TH sequence, each segment would have its own header. However, if each table segment after the first uses .TS H/.TH N then the table header will only appear at the beginning of the table and the top of each new page or column that the table continues onto.

7.4 Equations

```
.DS
.EQ [label]
equation(s)
.EN
.DE
```

The equation setters *eqn*(1) and *neqn*(1) [6,7] expect to use the .EQ (equation start) and .EN (equation end) macros as delimiters in the same way that *tbl*(1) uses .TS and .TE; however, .EQ and .EN must occur inside a .DS-.DE pair.

The .EQ macro takes an argument that will be used as a label for the equation. The label will appear at the right margin in the "vertical center" of the general equation.

☞ *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, Equation, and Exhibit Captions

```
.FG   [title]  [override]  [flag]
.TB   [title]  [override]  [flag]
.EC   [title]  [override]  [flag]
.EX   [title]  [override]  [flag]
```

The .FG (Figure Title), .TB (Table Title), .EC (Equation Caption) and .EX (Exhibit Caption) macros are normally used inside .DS-.DE pairs to automatically number and title figures, tables, and equations. They use registers *Fg, Tb, Ec,* and *Ex,* respectively.[14] 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", and .EX replaces "Figure" by "Exhibit". 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.

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, equation, and exhibit captions usually occur after the corresponding figures and equations.

7.6 List of Figures, Tables, Equations, and Exhibits

A List of Figures, List of Tables, List of Exhibits, and List of Equations may be obtained. They will be printed after the Table of Contents is printed if the number registers *Lf, Lt, Lx,* and *Le* (respectively) are set to 1. *Lf, Lt,* and *Lx* are 1 by default; *Le* is 0 by default.

7.7 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*(1) program may be used:

---

14. The user may wish to reset these registers after each first-level heading {4.6}.

```
  .
  .
  .
.DS 0 1
.TS
center;
lw40 .
T{
  .
  .
  .
T}
.TE
.DE
  .
  .
  .
```

The ".DS 0 1" begins a non-indented, filled display. The *tbl*(1) 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,[15] 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:

1. Automatically-numbered footnote:

   ```
   This is the line containing the word\*F
   .FS
   This is the text of the footnote.
   .FE
   to be footnoted.
   ```

---

15. Footnotes are processed in an environment that is different from that of the body of the text (see the .ev request in [9]).

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 unpaddable 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).

In a memorandum or a released paper, the page header on the first page is automatically 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 -rC*n* 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 *H1*):

    .PF ""'- \\\\n(H1-\\\\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 "Revision 3" on the top inside margin of each page:

```
.PH ""
.PF ""
.EH "'\\\\nP''Revision 3'"
.OH "'Revision 3''\\\\nP'"
```

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

```
.BS
zero or more lines of text
.BE
```

Lines of text that are specified between the .BS (bottom-block start) and .BE (bottom-block end) macros will be printed at the bottom of each page, after the footnotes (if any), but before the page footer. This block of text is removed by specifying an empty block, i.e.:

.BS
.BE

10. TABLE OF CONTENTS AND COVER SHEET

The table of contents and the cover sheet for a document are produced by invoking the .TC and .CS macros, respectively. The appropriate -rB*n* 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 comma. '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:

<div align="center">

Special Application
Message Transmission

Approved: _____

</div>

<div align="center">

CONTENTS

.

.

.

</div>

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.[16] 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*.[17] 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:

---

16. But only if -rB2 or -rB3 has been specified on the command line.
17. For ease of explanation, in this section (11.1) *nroff* behavior is described first, the convention of (1.2) not withstanding.

```
.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]

The .SA macro 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 *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 15.103 of the macros.

This information is useful in analyzing suspected bugs in PWB/MM. The easiest way to have this number appear in your output 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)
.1C
```

The .2C macro begins two-column processing which continues until a .1C 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. The .2C macro does *not* "balance" two-column output.

It is possible to have full-page width footnotes and displays when in two column mode, although the default action is for footnotes and displays to be narrow in two column mode and wide in one column

mode. Footnote and display width is controlled by a macro, .WC (Width Control), which takes the following arguments:

N	Normal default mode (— WF, — FF, — WD)
WF	Wide Footnotes always (even in two column mode)
—WF	DEFAULT: turn off WF (footnotes follow column mode, wide in 1C mode, narrow in 2C mode, unless FF is set)
FF	First Footnote; all footnotes have the same width as the *first* footnote encountered for that page
—FF	DEFAULT: turn off FF (footnote style follows the settings of WF or — WF)
WD	Wide Displays always (even in two column mode)
—WD	DEFAULT: Displays follow whichever column mode is in effect when the display is encountered

For example: .WC WD FF will cause all displays to be wide, and all footnotes on a page to be the same width, while .WC N will reinstate the default actions. If conflicting settings are given to .WC the last one is used. That is, .WC WF —WF has the effect of .WC —WF. Note that by default all options are turned off.

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 FORCING AN ODD PAGE

    .OP

This macro is used to ensure that the following text begins at the top of an odd-numbered page. If currently at the top of an odd page, no motion takes place. If currently on an even page, text resumes printing at the top of the next page. If currently on an odd page (but not at the top of the page) one blank page is produced, and printing resumes on the page after that.

11.9 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.[18]

12. ERRORS AND DEBUGGING

12.1 Error Terminations

When a macro discovers an error, the following actions occur:

- A break occurs.

---

18. 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.

- 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. (All the error messages are explained in Appendix E.)

- 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 300(1), 450(1), or hp(1) 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(1) or eqn(1)/neqn(1), 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.*

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: .S)
            % (page number)

13.1.2 Names Used by PWB/MM.

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 EQN/NEQN and TBL. The equation preprocessors, *eqn*(1) and *neqn*(1), use registers *and* string names of the form *nn.* The table preprocessor, *tbl*(1), uses names of the form:

```
    a-    a+    a|    nn    #a    ##    #-    #^    ^a    T&    TW
```

13.1.4 User-Definable Names. After the ̃·ove, what is left for user extensions? To avoid problems, we suggest using names that consist eitner 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 way of generating and numbering appendices: -

```
    .nr Hu 1
    .nr a 0
    .de aH
    .nr a +1
    .nr P 0
    .PH "'''Appendix \\na - \\\\\\\\\\nP'"
    .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.[19] 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:

```
.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.

---

19. The two-character sequence "\&" is understood by the formatters to be a "zero-width" space, i.e., it causes no output characters to appear.

- *Ease of use for simple documents*—It is not necessary to write complex sequences of commands to produce simple documents. 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.

- *Decoupling of input format from output style*—There is but one way to prepare the input text, although the user may obtain a number of output styles by setting a few global flags. For example, the .H macro is used for all numbered headings, yet the actual output style of these headings may be made to vary from document to document or, for that matter, within a single document.

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. Many of the features of PWB/MM are patterned after similar features in a number of earlier macro packages, and, in particular, after one implemented by M. E. Lesk. 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. Unpublished 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):[20]

```
.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
..
```

---

20. 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.

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:[21]

```
.de BL
.LB 3 0 2 0 \\•(BU
..
```

---

21. 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(.S .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\} \" and reset g to current level (.aL changes g)
.if \\n(:g>\\ng .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 exarⁿle, 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 footnote[1] 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 "May 31, 1979"
.TL 334455
Out-of-Hours Course Description
.AU "D. W. Stevenson" DWS PY 9876 5432 1X-123
.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
```

subject: Out-of-Hours Course Description   date: May 31, 1979
       Case: 334455

from: D. W. Stevenson
       PY 9876
       ıX-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 tнe 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: **May 31, 1979**

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:[22]

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 , 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: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,

---

22. This list is set up by ".LB 37 0 2 0" (5.4).

the message is issued because there almost certainly exists some problem in the preceding text.

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 *value*.

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 do ev" is caused by *(a)* setting a page width that is negative or extremely short, *(b)* setting a page length that is negative or extremely short, *(c)* reprocessing a macro package (e.g. performing a .so to a macro package that was requested from the command line), and *(d)* requesting the − sl option to *troff* on a document that is longer than ten pages.

"Cannot open *filename*" is issued if one of the files in the list of files to be processed cannot be opened.

"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.

"Non-existent font type" means that a request has been made to mount an unknown font.

"Non-existent macro file" means that the requested macro package does not exist.

"Non-existent terminal type" means that the terminal options refers to an unknown terminal type.

"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*(1)/*neqn*(1).

Appendix F: SUMMARY OF MACROS, STRINGS, AND NUMBER REGISTERS

L. 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 (6.3)
        .AU name [initials] [loc] [dept] [ext] [room] [arg] [arg] [arg]

B   Bold (underline in *nroff*) (11.1)
        .B [bold-arg] [previous-font-arg]

BE   Bottom End (9.13)
        .BE

BL   Bullet list start (5.3.3.2)
        .BL [text-indent] [1]

BS.   Bottom Start (9.13)
        .BS

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 [label]

EX   Exhibit caption {7.5}
.EX [title] [override] [flag]

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 *nroff*) {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] ...

OP   Odd page {11.7}
      .OP

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

TH   Table header (7.3)
.TH [N]

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 [H]

TX *   Table-of-contents user exit (10.1)
.TX

VL   Variable-item list start (5.3.3.6)
.VL text-indent [mark-indent] [1]

WC   Width Control (11.4)
.WC [format]

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., January 22, 1980)

EM   Em dash string, produces an em dash for both *nroff* and *troff* (3.8).

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*)

HP   Heading point size list, up to seven codes for heading levels 1 through 7 (4.2.2)

RE   SCCS Release and Level of PWB/MM (11.3)
Release.Level (e.g., 15.103)

Tm   Trademark string places the letters "TM" one half-line above the text that it follows.

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]

De   Display eject register for floating dislays (7.2)
        0, [0:1]

Df   Display format register for floating displays (7.2)
        5, [0:5]

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*)[23]

Li   List indent {5.3.3.1}
      5, [0:?]

Ls   List spacing between items by level {5.3.3.1}
      5, [0:5]

N •   Numbering style {2.4}
      0, [0:3]

O •   Offset of page {2.4}
      0, [0:?] (0.5i, [0i:?] in *troff*)[23]

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*)[23]

January 1980

---

23. For *nroff*, these values are *unscaled* numbers representing lines or character positions; for *troff*, these values must be *scaled*.
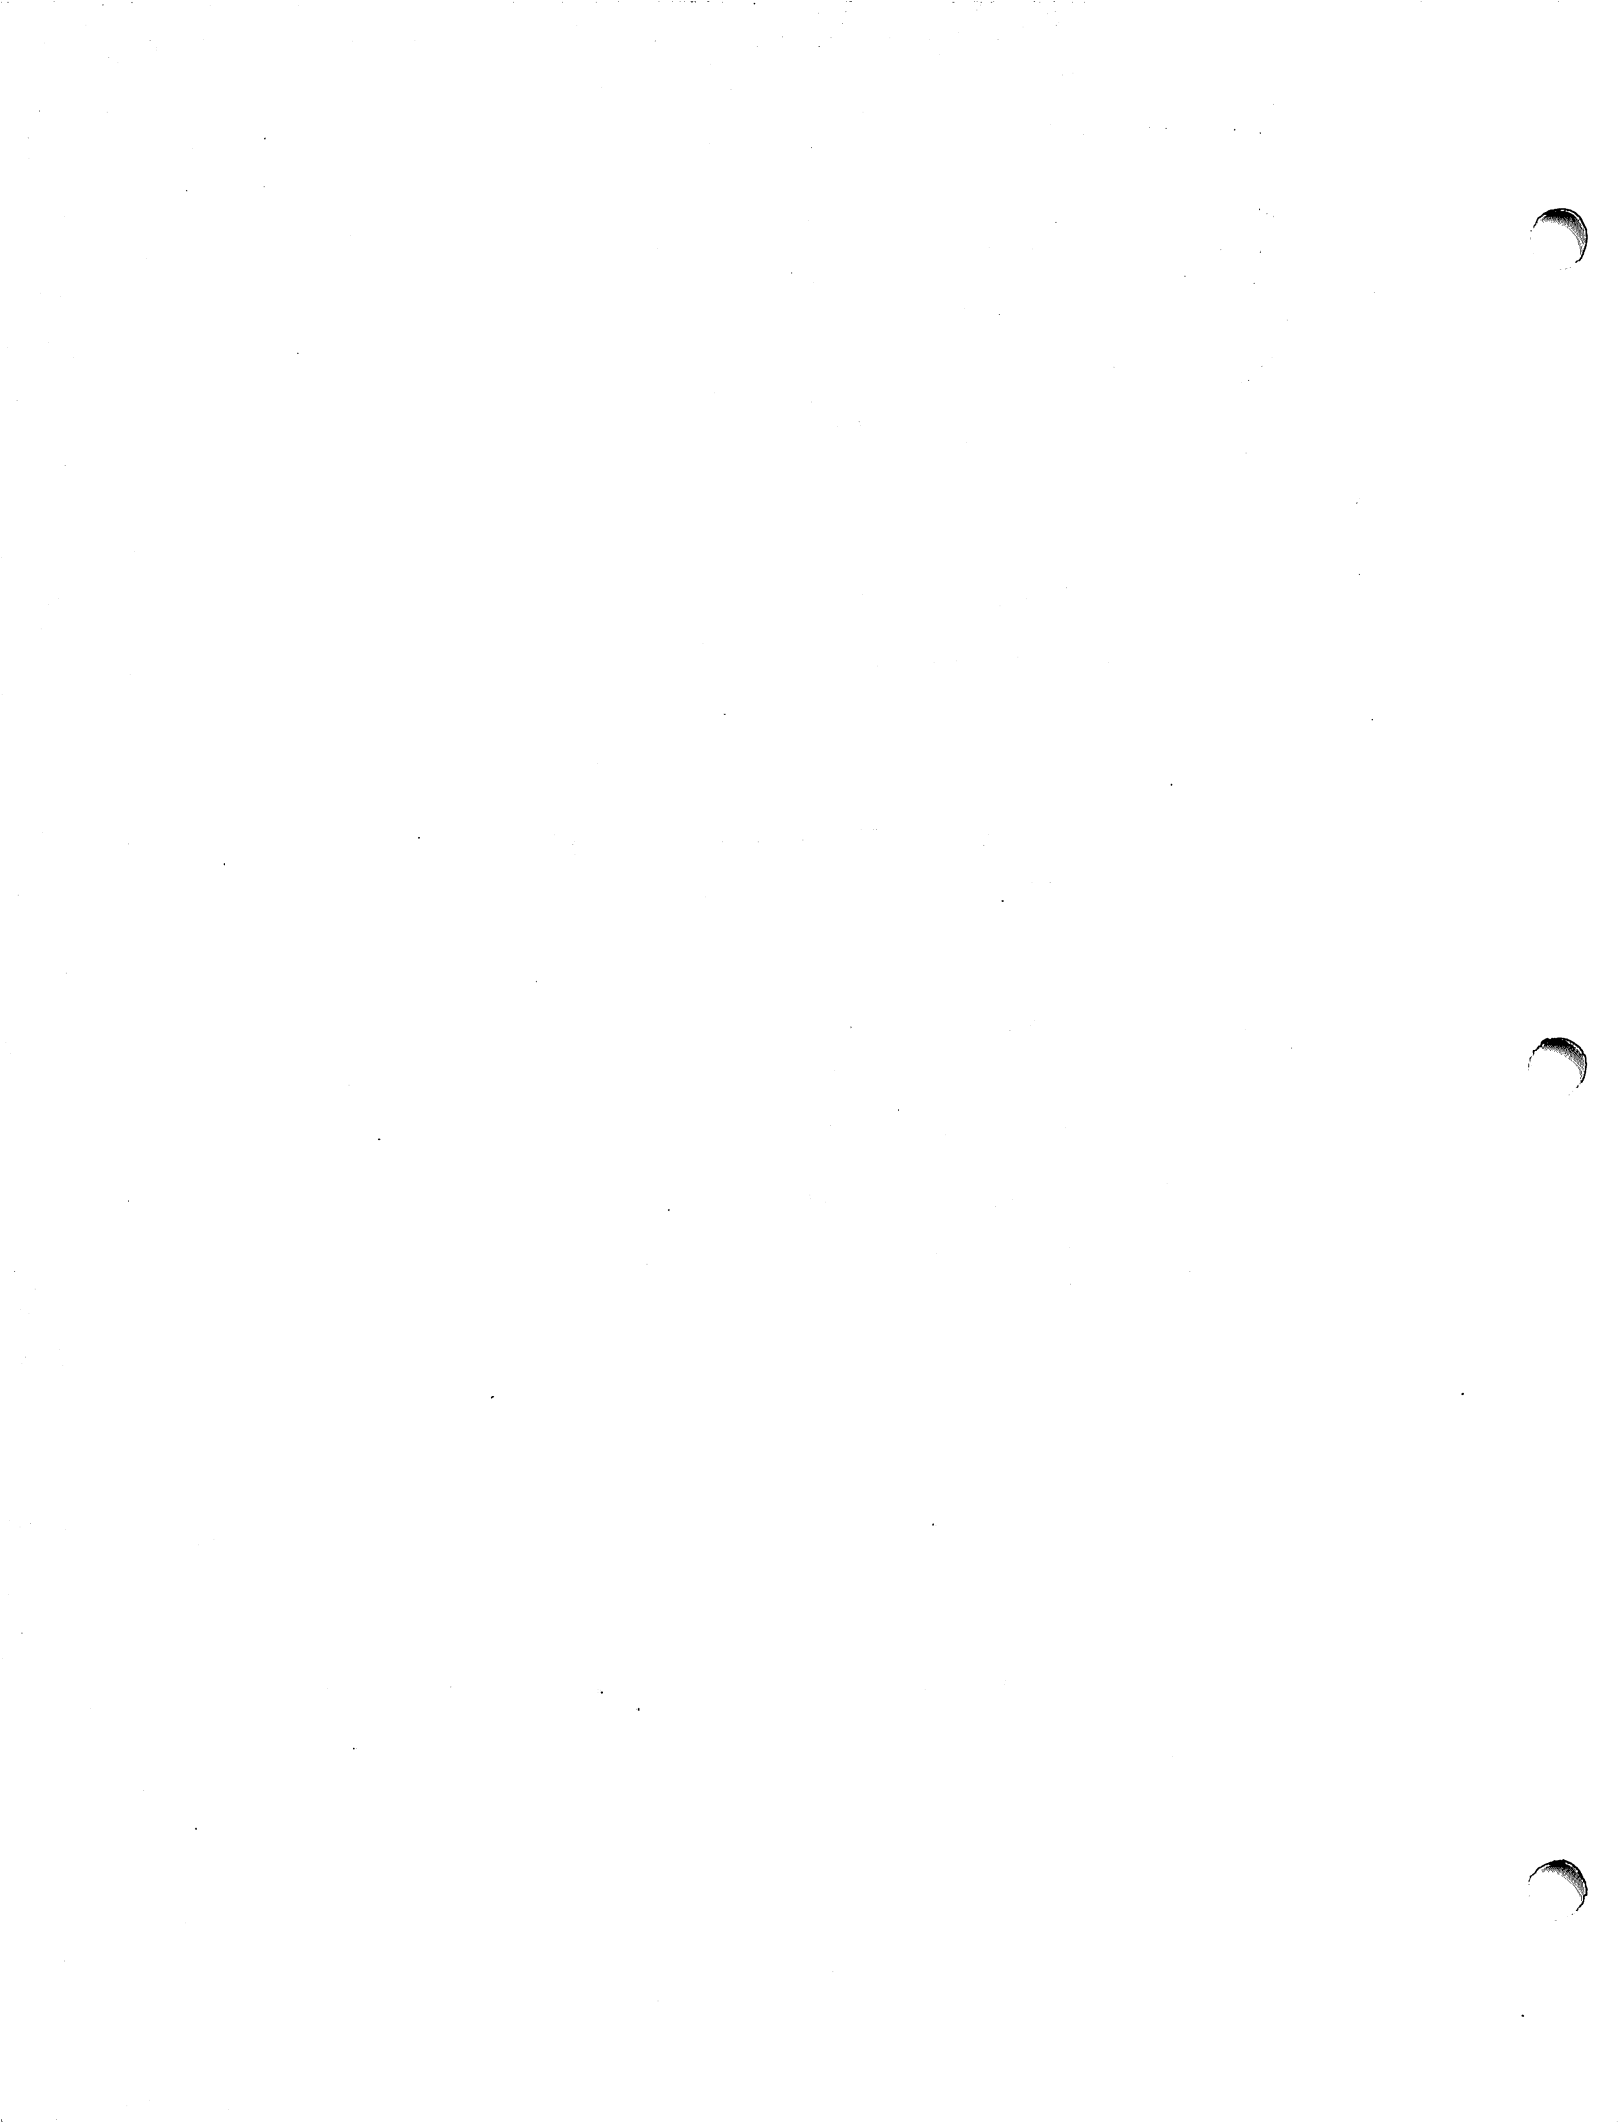
ADDENDUM

The following changes to the Memorandum Macros package are effective with the UNIX[TM] System III release.

Changes to the Memorandum Macros package fall into and are described under five categories: enhancements, bug fixes, documentation only changes, BTL Office Guide conformity, and the **mm** command.

1. MM ENHANCEMENTS

New Point Size and Vertical Spacing Macro (.S)

> .S [point size] [vertical spacing]

.S has been rewritten to accommodate a second argument for vertical spacing and to offer more flexibility to the user. The first argument still specifies point size. Mnemonics, such as D for default value, C for current value, and P for previous value, are now available and may be used for both point size and vertical spacing arguments.

Arguments may be signed or unsigned. If an argument is negative, the current value is decremented by the specified amount. If the argument is positive, the current value is incremented by the specified amount. If an argument is unsigned, it is used as the new value. .S without arguments defaults to previous (.P). If the first argument is specified but the second argument (vertical spacing) is not then the default (D) value is used. The default value for vertical spacing is always 2p greater than the current point size value selected. A null ("") argument for either the first or second argument defaults to the current (C) value.

Example: (where n is a numeric value):

```
.S              ->    .S P P
.S "" n         ->    .S C n
.S n ""         ->    .S n C
.S n            ->    .S n D
.S ""           ->    .S C D
.S "" ""        ->    .S C C
.S n n          ->    .S n n
```

By default, PWB/MM initializes point size to 10 and vertical spacing to 12.

If a point size argument is greater than 99, the default point size (D) 10 is restored. If a vertical spacing argument is greater than 99, the default vertical spacing (D) +2p is used.

EXAMPLE:

```
.S 12 111       ->    .S 12 14
.S 110          ->,   .S 10 12
```

Note: Remember, the default vertical spacing setting (D) is dependent upon the current point size.

New User-exit Macro .TY

```
.de TY
1 or more lines of text
..
```

By defining .TY, whose usage is identical to the user-exit macro .TX, the user may replace the word "CONTENTS" on the Table of Contents page.

EXAMPLE:

```
.de TY
..
```

Any text that the user wishes to print may be included in the definition as with .TX.

Memorandum for File Cover Sheet with Abstract and Copy to

```
.AS [type] [indent]
```

A Memorandum for File cover sheet with abstract and copy to list may be obtained by specifying .AS 2. This produces two pages with the usual memorandum header information, such as subject, date, from, etc. The first page includes the

abstract and copy to list. The second page includes the
memorandum type, such as Memorandum for File, Engineer's
Notes, etc, followed by the text.

Notations (.NS/.NE) may appear within the Memorandum for
File cover sheet. The sequence of calls is as follows:

EXAMPLE:
```
                              .TL
                              .AU
                              .AS 2
                              .AE
                              .NS
                              .NE
                              .MT
```

Notations are also allowed in the usual memorandum closing
location after the .SG. However, notations included on the
cover sheet are not repeated. If .AS 2 is not given, any
.NS/.NE's before the .MT are saved and output at the end of
the document.

If your memorandum requires that you assign a permuted index
number (sequence number), or some other additional
information, arguments 7 through 9 on the author line may be
used.

EXAMPLE:

.AU "N. E. Bock" NEB ... 2F-203 3646-800407.01MF

produces the from block on this memorandum.

New Options for .I/.B and New Font Change Macros'

```
    .I [arg1] ... [arg6]
    .B [arg1] ... [arg6]
```

.I and .B now take up to six arguments. Arguments 1, 3, and
5 are italic or bold for .I and .B respectively. Arguments
2, 4, and 6 revert to the previous font.

In addition, 6 new font change macros have been added (these
macros are the same as in the UNIX User's Manual Entry Macro
package).

                .IB, .BI, .RI, .IR, .RB, .BR

where I stands for Italic, B for Bold, and R for Roman.

These macros take up to six arguments. Arguments 1, 3, and
5 are produced in the font specified by the first letter of
the macro name, and arguments 2, 4, and 6 are produced in
the font specified by the second letter of the macro name.

In troff, .I, .B, .BI, .IB, .RI, and .IR separate their
arguments with a 1/12 space when one of the alternated fonts
is italic.

New Option for Displays

        .DS [format] [fill] [right indent]

A third argument has been added to displays (.DS/.DF) to
decrease the display line length by the amount of the
argument. The argument must be scaled in troff, otherwise
the default is em's, and the argument must be unscaled in
nroff, defaulting to en's.

This argument produces a right indent. The user must
specify the left indent.

EXAMPLE:
                        .DS L F 5

creates a display which is indented only from the right,
however,

                        .DS I F 5

creates a display which is indented from both sides.

New String for Indentation of Headings in Table of Contents

        .ds Ci indent

The string Ci (Contents indent) may be used to specify the
indentation of heading levels in the Table of Contents. The
Ci string takes seven arguments which must be scaled.

EXAMPLE:
                .ds Ci .25i .5i .75i 1i 1i

        or

        .ds Ci 0 2n 4n 6n 8n

Each level allowed by the Cl register must be specified in
this string. For the above example, Cl must not be set to a
value greater than 5.

New Register for Page Numbering of Table of Contents

.nr Oc value

The Oc number register allows the user to specify lower-case Roman or Arabic page numbering on the Table of Contents.

Value	Meaning
.nr Oc 0	lower-case Roman (default)
.nr Oc 1	Arabic

New Register for Placement of List of Figures, Tables, Equations, and Exhibits

.nr Cp value

The Cp number register allows the list of figures, tables, equations, and exhibits to be on separate pages following the Table of Contents or to be listed on the Table of Contents page.

Value	Meaning
.nr Cp 0	separate pages following Table of Contents (default)
.nr Cp 1	Table of Contents page

New Strings for Figure, Table, Equation, and Exhibit Titles for Table of Contents

.ds Lf [title]

The user may define strings to be the titles for these sections. The default settings for these strings are as follows:

```
.ds Lf LIST OF FIGURES
.ds Lt LIST OF TABLES
.ds Lx LIST OF EXHIBITS
.ds Le LIST OF EQUATIONS
```

New Register for Numbered Paragraphs

.nr Np value

Numbered paragraphs may be obtained by setting the Np number register.

Value	Meaning
.nr Np 0	unnumbered paragraphs (default)
.nr Np 1	numbered paragraphs

If numbered paragraphs are desired then each paragraph (.P) is numbered within first level heading sections, as follows: 1.01, 1.02, 1.03, 2.01, etc.

The P1 (one) number register is available to reset numbered paragraphs. By default, the P1 number register is reset to 0 by the H1 counter.

New Options for Width Control Macro (.WC)

        .WC code

The Width Control macro (.WC) has been expanded to control formatting of floating displays with two new options.

Code	Meaning
FB	floater break (default)
-FB	no floater break

These can be used to control breaks caused by floating displays (.DF) when output on the "current page". By default, floating displays cause a break when output on the current page.

New Accent Strings

New strings for letter accents have been added to PWB/MM.

	Input	Output
Grave accent	a*`	à
Acute accent	a*'	á
Circumflex	a*^	â
Tilde	n*~	ñ
Cedilla	c*,	ç
Lower-case umlaut	a*:	ä
Upper-case umlaut	A*;	Ä

Calculation of Cover Sheet Macro Arguments

        .CS [pages] [other] [total] [figures] [tables] [references]

The cover sheet macro (.CS) now calculates its arguments as
follows:

argument	register value
pages	\nP
other	0
total	\nP + other
figures	\n(Fg
tables	\n(Tb
references	\n(Rf

This automatic calculation may be obtained by calling .CS
with no arguments. However, if values need to be over-
ridden, i.e., second argument "other", then null arguments
may be used.

EXAMPLE:
                                .CS "" 3

Null arguments default to values shown in above table.

New Location Codes

        .AU name [initials] [location] . . . [arg9]

HOH (Crawford Hill) and WB (West Long Branch) are now
available as location codes for Memorandum for File and
Released Papers. HOH provides an extra line with "Crawford
Hill Laboratory" on Released Paper format.

New Register for Dot/Dash Separator in Figure, Table,
Equation, and Exhibit Captions

        .nr Of value

The Of number register allows a period or a dash to be used
in separating the figure, table, exhibit and equation
captions from titles.

Value	Meaning
.nr Of 0	Figure 1. Title (default)
.nr Of 1	Figure 1 - Title

New Register for Right/Left Adjusted Equation Labels

    .nr Eq value

The Eq number register is available to produce left or right
adjusted equation labels.

Value	Meaning
.nr Eq 0	right adjusted equation labels (default)
.nr Eq 1	left adjusted equation labels

Note: Left blocked (.DS) and left indented (.DS 1) equations
      are right adjusted when equation labels are left
      adjusted (.nr Eq 1).

Alternate Format Page Numbering

The output format of the ℗ number register may now be
changed via the .af (alternate format) request.

New Macro for Proprietary Marking (.PM)

    .PM [code]

.PM, for Proprietary Marking, appends to the page footer a
PRIVATE, NOTICE, BELL LABORATORIES PROPRIETARY, or BELL
LABORATORIES RESTRICTED disclaimer.   .PM takes one argument.

Code	Meaning
.PM	no proprietary marking
.PM P	PRIVATE disclaimer
.PM N	NOTICE disclaimer
.PM BP	BELL LABORATORIES PROPRIETARY disclaimer
.PM BR	BELL LABORATORIES RESTRICTED disclaimer

The disclaimers are in a form approved for use by the Bell
System.

New Stop Code Macro (.RD)

    .RD [prompt] [diversion] [string]

.RD (ReaD insertion) allows a user to stop the standard
output of a document and to read text from the standard
input until 2 consecutive newlines are found. When the
newlines are encountered, normal output is resumed. .RD may
take 3 arguments.

Note that .RD follows the formatting conventions in effect.
For example, the examples below assume that the .RD is
invoked in no fill mode (.nf).

The first argument is a prompt.

EXAMPLE:

                               .RD Name

.RD prints the text "Name:" and then stops the standard
output. The user may now fill-in this field. When 2
consecutive newlines are returned, the standard output
resumes, until another .RD is encountered or the job
terminates.

.RD without a prompt (1st argument) signals the user with a
BEL on terminal output.

The second argument, a diversion name, allows the user to
save all the entered text typed after the prompt.

EXAMPLE:
                               .RD Name aa bb
produces

            Name: (user types) J. Jones
            16 Elm Rd.,
            Piscataway

The diversion aa (.aa) would contain:

                    J. Jones
                    16 Elm Rd.,
                    Piscataway

The third argument, a string name, allows the user to save
for later reference the first line following the prompt.

EXAMPLE:
                               .RD Name "" bb
produces
            Name: (user types) J. Jones

In both examples, the string bb (*(bb) contains "J. Jones".

Note: A newline followed by a Control D (EOF) also allows
      the user to resume normal output.

New Macro for Vertical Margins (.VM)

    .VM [top] [bottom]

.VM (vertical margin) allows the user to specify extra space
at the top and bottom of the page. This space precedes the
page header and follows the page footer. .VM takes two
arguments. These arguments are scaled in v's.

EXAMPLE:
                              .VM 10 15

adds 10 blank lines to the default top of page margin, and
15 blank lines to the default bottom of page margin. Both
arguments must be positive (default spacing at the top of
the page may be changed by re-defining TP).

New Usage of Notation Blocks

    .NS [arg]
    zero or more lines of notation
    .NE

Notation blocks (.NS/.NE) may be used in documents which do
not contain an .MT. However, a .br must precede the .NS/.NE
in this case. All notations are saved and output at the end
of the document.

Resetting of Tabs by Center Blocked Displays

    .MT
    .ta tab1 tab2 ...

Center blocked displays (.DS CB) no longer reset tabs.
Thus, the user may set tabs for the entire document after
specifying the memorandum type (.MT).

New Double-line Indent Macro (.nP)

Double indented paragraphs are produced using .nP.
Paragraphs are numbered within second level headings.

EXAMPLE:
                        .H 1 first
                        .H 2 second
                        .nP
                        text of paragraph

Macro Call Only for Generation of Table of Contents and
Cover Sheet

    mm [options] file

The user no longer needs to specify the command line option
-rB[123] to produce a Table of Contents and/or Cover Sheet.
However, the macros .TC and .CS must be used.

New Command Line Option -rE1

    mm -rE1 [options] file

In mm (nroff), the bold printing of subject, date and from
fields and the signer's name may be turned on by using the
command line option -rE1. By default in mm, these fields
are not bold. In mmt (troff), the bold font is on by
default and may be turned off with -rE0.

Note: The company name is always emboldened.

Command Line Options to Change Page Width and Page Offset on
Preprinted Forms

    mm -rA1 -rWn -rOn [options] file

Page Width (5.8 inches) and Page Offset (1 inch) of
preprinted forms may be changed by using the command line
option -rA1, with -rWn (page width) and/or -rOn (page
offset), rather than by using a .AF.

New Command Line Option -rN5

    mm -rN5 [options] file

The -rN5 command line option produces section/page numbering
like -rN3 and also automatically numbers figure, table,
equation, and exhibit captions within first level heading
sections.

New Command Line Option -rA2

The -rA2 command line option allows for usage of the Bell
Logo, if available, on a printing device (currently
available for the Xerox 9700 only).

2. MM BUG FIXES

Fill/No Fill Mode Preserved within Displays

Figure, table, equation, and exhibit captions (.FG, .TB,
.EQ, and .EX respectively) preserve the fill/no fill mode
within displays. For example, a long figure caption in a no
fill display (.DS "" N) turns on fill mode for the caption,
but then restores no fill mode for the display.

Alignment of Captions within Displays

Long figure, table, equation, and exhibit captions now align
correctly inside of centered displays (.DS C).

Captions Centered with Respect to Left Margin

Figure, table, equation, and exhibit captions are now
correctly centered with respect to left margin rather than
to indentations.

Fill/No Fill Mode Restored after Displays

Displays restore fill/no fill mode in the text following the
display to the state preceding the display.

Point Size on MFF and Released Papers

In Memorandum for File and Released Paper format, the point
size is changed back to the previous (not default) after the
title (for Released Papers only) and the .MT.

List Indent

The Li number register (list indent) works as documented.

Floating Displays Forced Out before Signature Line

Any floating display not yet printed is forced out before
the signature line in a Memorandum for File.

Page Width and Offset Restored for Table of Contents

If a .TC (Table of Contents) follows a .CS (Cover Sheet),
the page width and offset are restored properly.

Indents Restored When Using Continuous Table Headers

Indents are no longer lost when using .TH N (continuous table header macro).

Labeled Footnotes on TM Cover Sheet

Bugs with labeled footnotes on a Technical Memorandum cover sheet have been fixed.

No Bottom of Page Processing in Cover Sheet

The .BS/.BE (bottom of page processing) block of text does not appear on the cover sheet for Technical Memorandum or Released Papers, although it still appears on the Table of Contents.

Centered Captions Outside of Displays

Figure, table, exhibit, and equation captions now center properly outside of displays.

Static Displays Made Consistent with Floating Displays

Static displays (.DS) are now output single spaced in order to be consistent with floating displays (.DF). However, line spacing (.ls) may be set inside of displays.

Cut Marks for Troff

Cut marks for _troff_ have been fixed to work properly (for the STARE).

Command Line Option -rN1

The command line option -rN1 now works properly with a one page Memorandum for File.

3. **MM DOCUMENTATION ONLY CHANGES**

The following are existing features which were previously undocumented or incorrectly documented.

Run-in Headings

Run-in headings which fit on one line are continuously underlined.

Hyphenation Indication

The hyphenation indication character may occur only at the beginning of a word.

Automatically Numbered Footnotes

Automatically numbered footnotes may not be used inside of tables.

Error Message for Page Length Specification

The error message "Cannot do ev" may be caused by a page length specification that is too short.

Naming Conventions for Strings, Macros, and Diversions

Warning: may not use the same name for strings, macros, and diversions within the same document.

Specification of Standard Input

A trailing minus on a nroff/troff command line is not needed to specify standard input if an input file is not named.

User-exit Macro .HY

The user-exit macro .HY is called shortly after the .ne request is done for the size of the heading text. The same arguments passed to .HX are passed to .HY. .HY may be used to restore temporary indents which may have been lost during processing.

User-exit Macro .HX

The string }0, used in the .HX, contains the heading mark. Two paddable spaces are appended to this string.

Third Argument to Heading Macro

The heading macro (.H) takes a third argument. This argument is appended to the second argument. It is primarily used for a footnote mark.

Multiple Invocation of -mm

Invoking -mm more than once results in a proper error message and the job terminates.

Non-standard BTL Location Codes and Addresses

An author's location and address, which is a non-standard
BTL location code, may be defined before the author line
(.AU).

EXAMPLE:

> .ds xx [location-string]
> .AU name initials xx

The string xx is a two character name used for the location
code. The "location-string" is a one line company address
which appears on Released Paper format.

4. MM BTL OFFICE GUIDE CONFORMITY

In order to conform with the BTL Office Guide the following
enhancements have been made to PWB/MM. The necessary
changes were determined in a series of meetings with the
UNIX Task Force of the Word Processing Subcommittee. For
further details of these changes and also a list of the
accompanying changes made to the BTL Office Guide, see
"Final Report on PWB/MM Conformance to the BTL Office
Guide", by Nancy W. Smith, November 26, 1979, Case 39382-
900.

New Approval Signature Macro (.AV)

> .AV approver's-name

.AV (approval signature) may be placed after the last
notation block to generate an approval signature line.

EXAMPLE:

> .SG
> .NS 3
> .NS 5
> .NE
> .AV "approver's name"

New Formal Closing Macro (.FC)

> .FC [formal closing]

.FC (formal closing) may be placed immediately before the
.SG. This produces "Yours very truly," as a formal closing
lined up with the signer's name. If a different closing is
required, specify that closing as the first argument to .FC.

Second Argument to .MT for Addressee Name

        .MT [type] [addressee name]

.MT may take a second argument for addressee's name. The
addressee's name appears left adjusted in combination with
the page number on the second and following pages.

EXAMPLE:
                        .MT 1 "John Jones"

produces:

John Jones - 2

Note: This feature may not be used with Released Paper
      format (.MT 4 1).

New Macro for Author Title (.AT)

        .AT author-title

.AT (author title) allows the user to specify an author's
title, (i.e., Head, Education Center) which appears below
the signer's name on the signature line.

.AT must be given after the .AU for that author. .AT takes
9 arguments, each argument appears on a separate line
following the signer's name.

EXAMPLE:
                .AU name
                .AT "Head, Education Center"

New Register for PRIVATE Header

        .nr Pv value

The Pv number register outputs the word PRIVATE, centered
and underlined on the second line of a document (preceding
the page header). This is done by setting the Pv register.

Value	Meaning
.nr Pv 0	do not print PRIVATE (default)
.nr Pv 1	PRIVATE on first page only
.nr Pv 2	PRIVATE on all pages

If Pv is 2, the user definable .TP may not be used as .TP is
used by PWB/MM to print PRIVATE on all pages, except on the
first page of a memorandum which does not use .TP.

New Reference Macros

```
A line of text to be referenced.\*(Rf
.RS [string name]
reference text
.RF
```

Automatically numbered references may be obtained by typing
*(Rf immediately after the text to be referenced. This
places the next sequential reference number enclosed in
brackets a half-line above the text to be referenced. _Troff_
produces the reference number in a smaller point size.

There are two macros used to delimit text for each
reference. .RS takes one argument, a string name.

EXAMPLE:
```
.RS AA
reference text
.RF
```

The string AA is assigned the current reference number. It
may be used later in the document, as the string call,
*(AA, to reference text which must be labeled with a prior
reference number. The reference is output enclosed in
brackets a half-line above the text to be referenced. No
.RS/.RF is needed for subsequent references.

An automatically generated reference page is produced at the
end of the document before the Table of Contents and the
Cover Sheet are output. The reference page is entitled
"References", and is centered and underlined. This page
contains the reference text (RS/RF). The user may change
the Reference Page title by defining the Reference page
string.

EXAMPLE:
```
.ds Rp "New Title"
```

The optional .RP (Reference Page) macro may be used to
produce reference pages anywhere within a document (i.e.,
within heading sections). .RP takes two arguments.

```
.RP [arg1] [arg2]
```

These arguments allow the user to control resetting of
reference numbering, and page skipping.

arg1	Meaning
0	reset reference counter (default)
1	do not reset reference counter
arg2	Meaning
0	cause a following .SK (default)
1	do not cause a following .SK

Note: Because a reference page is produced automatically at
      the end of the document, .RP need not be used unless
      the user wishes to produce reference pages elsewhere
      in the document.

Page Header

Three blank lines are printed before and after the page
header. Previously two blank lines were printed.

Footnote Rule

The footnote rule is one inch long, preceded by at least two
blank lines, and followed by one blank line.

TM Number Positioning

The TM line on a technical memorandum is now correctly
aligned with the text of the "date" and "from" fields.

Signature Positioning

The signer's name on memoranda is now printed beginning at
the center of the page.

Charging and Filing Cases

Charging and Filing Cases on memoranda are handled as
follows:

   a. If there is a single number for Charging or Filing
       Case, the number immediately follows the text of the
       subject separated by a space hyphen space (" - ") and
       the word "Case".

   b. If the Charging and Filing Cases have the same number,
       "Charging and Filing Case" followed by the number
       appears on a separate line.

   c. If there are different Charging and Filing Cases,
       "Charging Case" followed by a number, and "File Case"
       followed by a number appear on separate lines.

Copy to

Continued notations, such as "Copy to", now appear as
follows:

                    Copy to - contd.

List Indent

The Li number register (list indent) defaults to 6 in _nroff_
and 5 in _troff_. Previously the values were 5 and 4 for
nroff and _troff_ respectively. The new indent provides two
spaces preceding a one digit number and one space preceding
a two digit number.

Hyphenation

Hyphenation is turned off by default.

Continuously Underlined Memorandum Type

Memorandum type (.MT) is now continuously underlined.

TM and Released Paper Abstracts

Abstracts no longer appear on page 1. Therefore,

                          .AS
                          .AS ""
                          .AS 0
                   and   .AS 1

are now equivalent. See PWB/MM ENHANCEMENTS, for new .AS 2
style.

Abstracts are no longer indented from the right and left
margins. To specify an abstract indent, the second argument
to .AS may be used.

EXAMPLE:
                      .AS "" 5

produces an abstract indented 5 spaces from the right and
left margins.

New Command Line Option for Draft Copies (-rC4)

    mm -rC4 [options] file

The command line option -rC4 for draft copies produces
double spacing, and indented paragraphs (with an indent of

10 spaces). The spacing takes effect after the .MT. (The command line option -rC3 remains the same, i.e., single spaced drafts.)

Page Offset

The default page offset in mm (nroff) is set to 3/4 inch. Since the average terminal prints approximately 1/2 inch from the left margin, the new page offset produces a left margin of 1 1/4 inches. The default page offset in mm (nroff) for pre-printed forms is 1 inch. This setting may be over-ridden by specifying -rOn on the command line (where n is characters and may equal 0).

Line Length

The line length now defaults to 6.0 inches (both mm and mmt), and on pre-printed forms (command line option -rA1 or .AF in text file) to 5.8 inches. (Pre-printed forms are mm only). .

5. MM COMMAND (mm,mmt)

New Default Terminal Type

By default, the mm shell command will use 450 as the value of term (-T450).

New Option for Non-compacted Macros (-y)

The -y option (replacing old usage of -u option) allows users to request non-compacted macros in the mm, mmt, and man shell procedures.

EXAMPLE:
                          mm -y [options] file

(Uses -mm rather than -cm.)

New Versatec Printer Option (-Tvp)

The -Tvp option directs the output to a Versatec printer via the vpr(1) spooler. This option is available at MH only and may be used with the mmt and man commands.

New Terminal Type Options

The following are new terminal type options for the mm and man shell procedures.

-T40/4	Teletype model 40/4
-Ttek, -T4014	Tektronix 4014
-T2640, -T2621, -T2645	hps
-T382	DTC 382
-T1620, -T1620-12	Diablo 1620 (alias -T450, -T450-12)
-T4000A	Trendata 4000A
-T31, -T42	ADM31 and ADM42
-T33, -T35	tty terminals
-T2631	HP2631A line printer
-T2631-e	HP2631A line printer, expanded
-T2631-c	HP2631A line printer, compressed
-TX	text train line printers
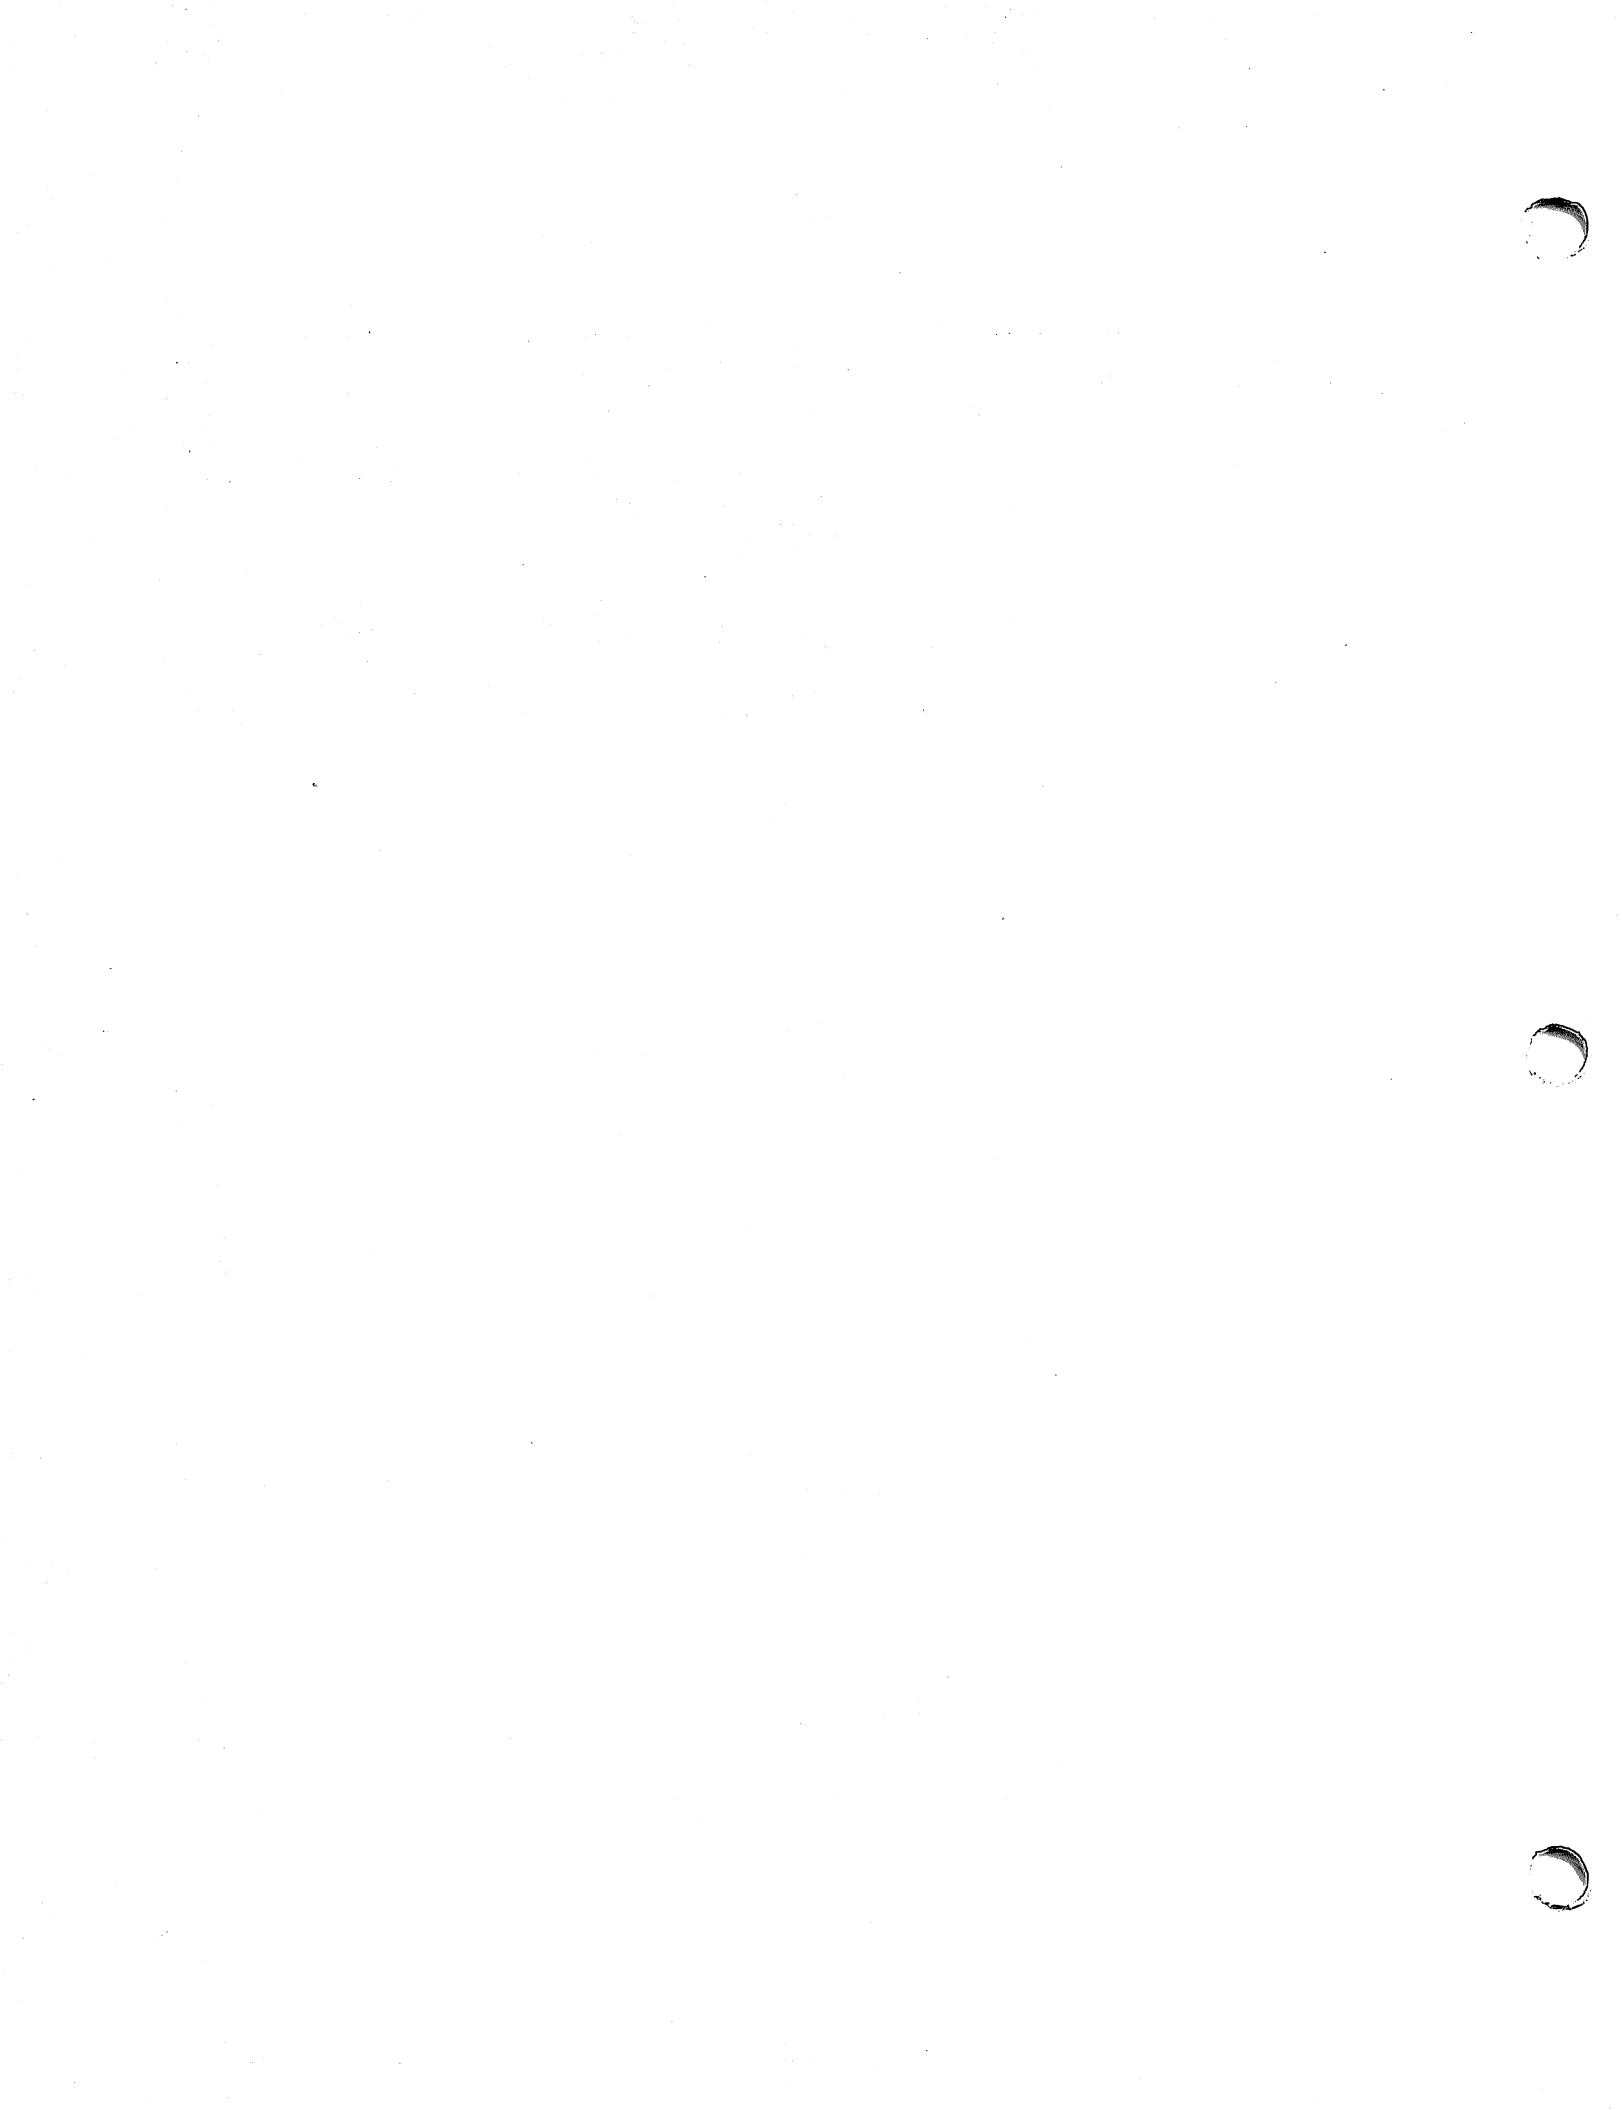
Tbl — A Program to Format Tables

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

January 16, 1979

---

* 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.

    linesize (n) — set lines or rules (e.g. from **box**) in n point type;

    delim (xy) — recognize x and y as the *eqn* delimiters.

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 l 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 an adjacent column contains a horizontal line, or if there are vertical lines adjoining this column, this horizontal line is extended to meet the nearby lines. 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

---

* More precisely, an en is a number of points (1 point = 1/72 inch) equal to half the current type size.

number is 3. If the separation is changed the worst case (largest space requested) governs.

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.

Vertical spacing changes — A key-letter may be followed by the letter **v** or **V** and a number to indicate the vertical line spacing to be used within a multi-line corresponding table entry. The number may be a signed digit, in which case it is taken as an increment or decrement from the current vertical spacing. A column separation value must be separated by blanks or some other specification from a vertical spacing request. This request has no effect unless the corresponding table entry is a text block (see below).

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)fI 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*. Such lines are extended to meet horizontal or vertical lines adjoining this 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.

Short horizontal lines — An input table *entry* containing only the string _ is taken to be a single line as wide as the contents of the column. It is not extended to meet adjoining lines.

Repeated characters — An input table *entry* containing only a string of the form \R*x* where *x* is any character is replaced by repetitions of the character *x* as wide as the data in the column. The sequence of *x*'s is not extended to meet adjoining columns.

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, spacing and font, using the p, v 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 10 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 directly.

For the convenience of users employing line printers without adequate driving tables or post-filters, there is a special −TX command line option to *tbl* which produces output that does not have fractional line motions in it. The only other command line options recognized by *tbl* are −ms and −mm which are turned into commands to fetch the corresponding macro files; usually it is more convenient to place these arguments on the *troff* part of the command line, but they are accepted by *tbl* as well.

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. The user can defend against this by giving the *delim(xx)* table option; this prevents splitting of numerical columns within the delimiters. For example, if the *eqn* delimiters are *$$*, giving *delim($$)* a numerical column such as "1245 $+- 16$" will be divided after 1245, not after 16.

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 ".TH" 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
    l l l.
    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 L̄.
january ⊕february ⊕march
april ⊕may
june ⊕july ⊕Months
august ⊕september
october ⊕november ⊕december
.TE
```

Output:

january	february	march
april	may	
june	july	**Months**
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;
cfI 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 formations; also Watchungs and Palisades.	200 million
Cenozoic	Coastal Plain	On Long Island 30,000 years; Cretaceous sediments redeposited by recent glaciation.

Input:

```
.EQ
delim $$
.EN

. . .

.TS
doublebox;
c c
l l.
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	$\sin(x)=\dfrac{1}{2i}(e^{ix}-e^{-ix})$
Error	$\mathrm{erf}(z)=\dfrac{2}{\sqrt{\pi}}\int_0^z e^{-t^2}dt$
Bessel	$J_0(z)=\dfrac{1}{\pi}\int_0^\pi \cos(z\sin\theta)\,d\theta$
Zeta	$\zeta(s)=\sum\limits_{k=1}^{\infty} k^{-s} \quad (\mathrm{Re}\ s>1)$

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
lB 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
lB 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 �crossAddress ⊕ 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) | ltw(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	
American Museum of Natural History	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
Bronx Zoo	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
Brooklyn Museum	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
New-York Historical Society	All the original paintings for Audubon's *Birds of America* 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.

References.

[1]   J. F. Ossanna, *NROFF/TROFF User's Manual,* Computing Science Technical Report No. 54, 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,* UNIX Programmer's Manual, Volume 2.

- 17 -

[5] M. E. Lesk and B. W. Kernighan, *Computer Typesetting of Technical Journals on UNIX, Proc. AFIPS NCC*, vol. 46, pp. 879-888 (1977).

[6] J. R. Mashey and D. W. Smith, "Documentation Tools and Techniques," *Proc. 2nd Int. Conf. on Software Engineering*, pp. 177-181 (October, 1976).

List of Tbl Command Characters and Words

Command	Meaning	Section
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
nnn	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
v V	Vertical spacing change	2
w W	Minimum width value	2
.*xx*	Included *troff* command	3
|	Vertical line	2
||	Double vertical line	2
^	Vertical span	2
\^	Vertical span	3
=	Double horizontal line	2,3
_	Horizontal line	2,3
_	Short horizontal line	3
\R*x*	Repeat character	3

Typesetting Mathematics — User's Guide   (Second Edition)

Brian W. Kernighan and 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 photo-typesetters 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 $\lim_{x \to \pi/2} (\tan x)^{\sin 2x} = 1$ or display equations like

$$G(z) = e^{\ln G(z)} = \exp\left(\sum_{k \geqslant 1} \frac{S_k z^k}{k}\right) = \prod_{k \geqslant 1} e^{S_k z^k / k}$$

$$= \left(1 + S_1 z + \frac{S_1^2 z^2}{2!} + \cdots\right)\left(1 + \frac{S_2 z^2}{2} + \frac{S_2^2 z^4}{2^2 \cdot 2!} + \cdots\right) \cdots$$

$$= \sum_{m \geqslant 0} \left(\sum_{\substack{k_1, k_2, \ldots, k_m \geqslant 0 \\ k_1 + 2k_2 + \cdots + mk_m = m}} \frac{S_1^{k_1}}{1^{k_1} k_1!} \frac{S_2^{k_2}}{2^{k_2} k_2!} \cdots \frac{S_m^{k_m}}{m^{k_m} k_m!} \right) z^m$$

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.

August 15, 1978

---

†UNIX is a Trademark of Bell Laboratories.

Typesetting Mathematics — User's Guide   (Second Edition)

Brian W. Kernighan and 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 comn.on 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 l. 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

$$x=f(y/2)+y/2 \qquad\qquad (3.1a)$$

There is also a shorthand notation so in-line expressions like π_i^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:

$$x~=~y~+~z$$

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,

$$x=2 \text{ pi int sin (omega t)dt}$$

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 \(bs for the Bell System sign Ⓐ.

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:

$$x~=~2~pi~int~sin~(~omega~t~)~dt$$

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 sub2* will give you *xsub2* 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 super-scripted 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:

$$e \text{ sup } \{i \text{ pi sup } \{rho +1\}\}$$

is

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 $\dfrac{-b^2}{\pi}$ 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:*

$$\text{sqrt } a+b + 1 \text{ over sqrt } \{ax \text{ sup } 2 +bx+c\}$$

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 over b sub 2}

is

$$\sqrt{\dfrac{a^2}{b_2}}$$

Big square roots are generally better written as something to the power ½:

$$(a^2/b_2)^{1/2}$$

which is

(a sup 2 /b sub 2) sup half

11. Summation, Integral, Etc.

Summations, integrals, and similar constructions are easy:

sum from i=0 to {i= 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,

$$\int \quad \prod \quad \bigcup \quad \bigcap$$

Since the thing before the *from* can be anything, even something in braces, *from-to* can often be used in unexpected ways:

lim from {n —> inf} x sub n =0

is

$$\lim_{n \to \infty} x_n = 0$$

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

bold x y

is

$$\mathbf{x}y$$

and

size 14 bold x = y +
    size 14 {alpha + beta}

gives

$$\mathbf{x} = y + \alpha + \beta$$

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

size 12 { ... }

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 thoughout 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.

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	$\overleftrightarrow{x}$
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 ^{2}He, 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 \ *(bs* 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

---

‡Like this one, in which we have a few random expressions like x_i and π^2. The sizes for these were set by the command *gsize −2*.

```
.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 .EQ I or .EQ L. 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

$$\left\{ \frac{a}{b}+1 \right\} = \left[\frac{c}{d} \right] + \left[e \right]$$

The resulting brackets are made big enough to cover whatever they enclose. Other characters can be used besides these, but the 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

$$\left\lfloor \frac{x}{y} \right\rfloor \leqslant \left\lceil \frac{a}{b} \right\rceil$$

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

$$x_i \quad x^2$$
$$y_i \quad y^2$$

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 $alpha sub i$ be the primary variable, and let $beta$ be zero. Then we can show that $x sub 1$ 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 sub i sub 1 + y sub i sub 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 xi1 ' 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.

22. A Large Example

Here is the complete source for the three display equations in the abstract of this guide.

```
.EQ I
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 I
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 I
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.

>=	$\geq$
<=	$\leq$
==	$\equiv$
!=	$\neq$
+-	$\pm$
->	$\rightarrow$
<-	$\leftarrow$
<<	$\ll$
>>	$\gg$
inf	∞
partial	∂
half	$\frac{1}{2}$
prime	$'$
approx	$\approx$
nothing	
cdot	$\cdot$
times	$\times$
del	∇
grad	∇
...	$\cdots$
,....,	$,\cdots,$
sum	$\sum$
int	$\int$
prod	$\prod$
union	$\cup$
inter	$\cap$

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	o
SIGMA	Σ	phi	ϕ
THETA	Θ	pi	π
UPSILON	Y	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
fwd	21	to	11
gfont	12	under	13
gsize	12	up	21
hat	13	vec	13
italic	12	~ ^	4, 6
lcol	18	{ }	8
left	16	"..."	8, 14
lineup	15		

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 .EQ ... .EN 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 —T*x*

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. 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, "NROFF/TROFF User's Manual", Bell Laboratories Computing Science Technical Report #54, 1976.

[2] M. E. Lesk, "Typing Documents on UNIX", Bell Laboratories, 1976.

[3] M. E. Lesk, "TBL — A Program for Setting Tables", Bell Laboratories Computing Science Technical Report #49, 1976.
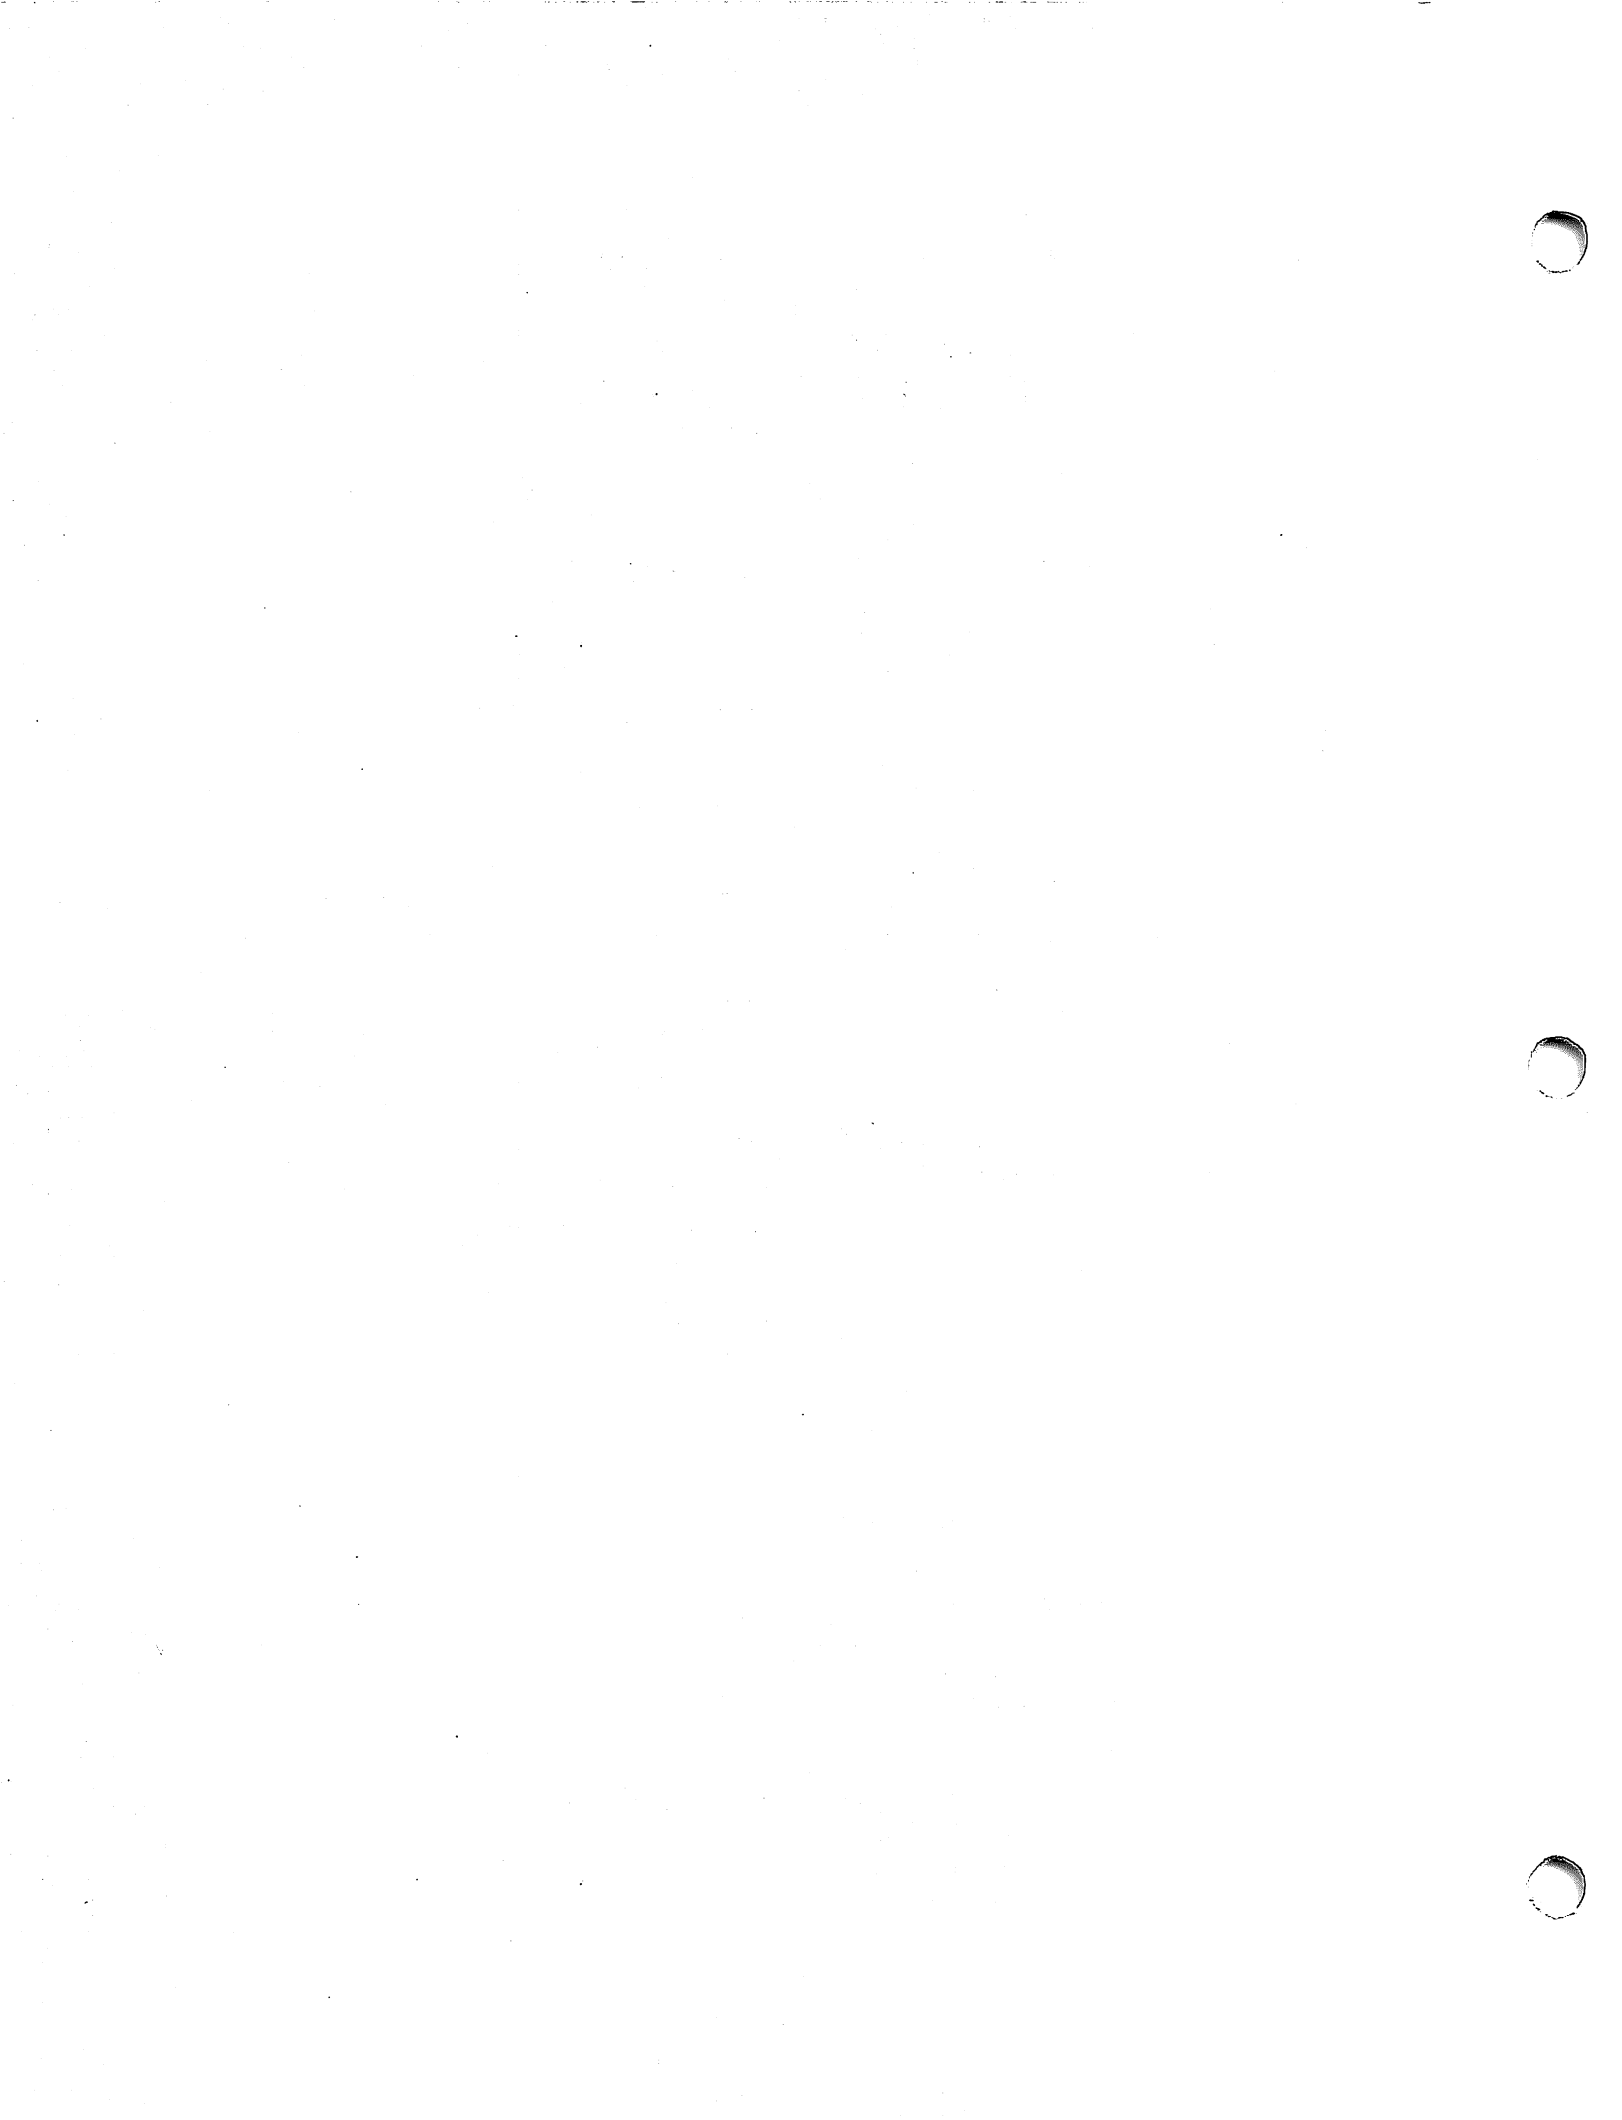
A System for Typesetting Mathematics

Brian W. Kernighan and Lorinda L. Cherry

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

This paper describes the design and implementation of a system for typesetting mathematics. The language has been designed to be easy to learn and to use by people (for example, secretaries and mathematical typists) who know neither mathematics nor typesetting. Experience indicates that the language can be learned in an hour or so, for it has few rules and fewer exceptions. For typical expressions, the size and font changes, positioning, line drawing, and the like necessary to print according to mathematical conventions are all done automatically. For example, the input

sum from i=0 to infinity x sub i = pi over 2

produces

$$\sum_{i=0}^{\infty} x_i = \frac{\pi}{2}$$

The syntax of the language is specified by a small context-free grammar; a compiler-compiler is used to make a compiler that translates this language into typesetting commands. Output may be produced on either a phototypesetter or on a terminal with forward and reverse half-line motions. The system interfaces directly with text formatting programs, so mixtures of text and mathematics may be handled simply.

This paper is a revision of a paper originally published in CACM, March, 1975.

1. Introduction

"Mathematics is known in the trade as *difficult*, or *penalty*, *copy* because it is slower, more difficult, and more expensive to set in type than any other kind of copy normally occurring in books and journals." [1]

One difficulty with mathematical text is the multiplicity of characters, sizes, and fonts. An expression such as

$$\lim_{x \to \pi/2} (\tan x)^{\sin 2x} = 1$$

requires an intimate mixture of roman, italic and greek letters, in three sizes, and a special character or two. ("Requires" is perhaps the wrong word, but mathematics has its own typographical conventions which are quite different from those of ordinary text.) Typesetting such an expression by traditional methods is still an essentially manual operation.

A second difficulty is the two dimensional character of mathematics, which the superscript and limits in the preceding example showed in its simplest form. This is carried further by

$$a_0 + \cfrac{b_1}{a_1 + \cfrac{b_2}{a_2 + \cfrac{b_3}{a_3 + \cdots}}}$$

and still further by

$$\int \frac{dx}{ae^{mx} - be^{-mx}} = \begin{cases} \dfrac{1}{2m\sqrt{ab}} \log \dfrac{\sqrt{a}\,e^{mx} - \sqrt{b}}{\sqrt{a}\,e^{mx} + \sqrt{b}} \\[2ex] \dfrac{1}{m\sqrt{ab}} \tanh^{-1}(\dfrac{\sqrt{a}}{\sqrt{b}} e^{mx}) \\[2ex] \dfrac{-1}{m\sqrt{ab}} \coth^{-1}(\dfrac{\sqrt{a}}{\sqrt{b}} e^{mx}) \end{cases}$$

These examples also show line-drawing, built-up characters like braces and radicals, and a spectrum of positioning problems. (Section 6 shows

what a user has to type to produce these on our system.)

2. Photocomposition

Photocomposition techniques can be used to solve some of the problems of typesetting mathematics. A phototypesetter is a device which exposes a piece of photographic paper or film, placing characters wherever they are wanted. The Graphic Systems phototypesetter[2] on the UNIX operating system[3] works by shining light through a character stencil. The character is made the right size by lenses, and the light beam directed by fiber optics to the desired place on a piece of photographic paper. The exposed paper is developed and typically used in some form of photo-offset reproduction.

On UNIX, the phototypesetter is driven by a formatting program called TROFF [4]. TROFF was designed for setting running text. It also provides all of the facilities that one needs for doing mathematics, such as arbitrary horizontal and vertical motions, line-drawing, size changing, but the syntax for describing these special operations is difficult to learn, and difficult even for experienced users to type correctly.

For this reason we decided to use TROFF as an "assembly language," by designing a language for describing mathematical expressions, and compiling it into TROFF.

3. Language Design

The fundamental principle upon which we based our language design is that the language should be easy to use by people (for example, secretaries) who know neither mathematics nor typesetting.

This principle implies several things. First, "normal" mathematical conventions about operator precedence, parentheses, and the like cannot be used, for to give special meaning to such characters means that the user has to understand what he or she is typing. Thus the language should not assume, for instance, that parentheses are always balanced, for they are not in the half-open interval $(a,b]$. Nor should it assume that that $\sqrt{a+b}$ can be replaced by $(a+b)^{1/2}$, or that $1/(1-x)$ is better written as $\frac{1}{1-x}$ (or vice versa).

Second, there should be relatively few rules, keywords, special symbols and operators, and the like. This keeps the language easy to learn and remember. Furthermore, there should be few exceptions to the rules that do exist: if something works in one situation, it should work everywhere. If a variable can have a subscript, then a subscript can have a subscript, and so on

without limit.

Third, "standard" things should happen automatically. Someone who types "x=y+z+1" should get "$x=y+z+1$". Subscripts and superscripts should automatically be printed in an appropriately smaller size, with no special intervention. Fraction bars have to be made the right length and positioned at the right height. And so on. Indeed a mechanism for overriding default actions has to exist, but its application is the exception, not the rule.

We assume that the typist has a reasonable picture (a two-dimensional representation) of the desired final form, as might be handwritten by the author of a paper. We also assume that the input is typed on a computer terminal much like an ordinary typewriter. This implies an input alphabet of perhaps 100 characters, none of them special.

A secondary, but still important, goal in our design was that the system should be easy to implement, since neither of the authors had any desire to make a long-term project of it. Since our design was not firm, it was also necessary that the program be easy to change at any time.

To make the program easy to build and to change, and to guarantee regularity ("it should work everywhere"), the language is defined by a context-free grammar, described in Section 5. The compiler for the language was built using a compiler-compiler.

A priori, the grammar/compiler-compiler approach seemed the right thing to do. Our subsequent experience leads us to believe that any other course would have been folly. The original language was designed in a few days. Construction of a working system sufficient to try significant examples required perhaps a person-month. Since then, we have spent a modest amount of additional time over several years tuning, adding facilities, and occasionally changing the language as users make criticisms and suggestions.

We also decided quite early that we would let TROFF do our work for us whenever possible. TROFF is quite a powerful program, with a macro facility, text and arithmetic variables, numerical computation and testing, and conditional branching. Thus we have been able to avoid writing a lot of mundane but tricky software. For example, we store no text strings, but simply pass them on to TROFF. Thus we avoid having to write a storage management package. Furthermore, we have been able to isolate ourselves from most details of the particular device and character set currently in use. For example, we let TROFF compute the widths of all strings of

characters; we need know nothing about them.

A third design goal is special to our environment. Since our program is only useful for typesetting mathematics, it is necessary that it interface cleanly with the underlying typesetting language for the benefit of users who want to set intermingled mathematics and text (the usual case). The standard mode of operation is that when a document is typed, mathematical expressions are input as part of the text, but marked by user settable delimiters. The program reads this input and treats as comments those things which are not mathematics, simply passing them through untouched. At the same time it converts the mathematical input into the necessary TROFF commands. The resulting ioutput is passed directly to TROFF where the comments and the mathematical parts both become text and/or TROFF commands.

4. The Language

We will not try to describe the language precisely here; interested readers may refer to the appendix for more details. Throughout this section, we will write expressions exactly as they are handed to the typesetting program (hereinafter called "EQN"), except that we won't show the delimiters that the user types to mark the beginning and end of the expression. The interface between EQN and TROFF is described at the end of this section.

As we said, typing $x=y+z+1$ should produce $x=y+z+1$, and indeed it does. Variables are made italic, operators and digits become roman, and normal spacings between letters and operators are altered slightly to give a more pleasing appearance.

Input is free-form. Spaces and new lines in the input are used by EQN to separate pieces of the input; they are not used to create space in the output. Thus

```
x    =    y
  + z + 1
```

also gives $x=y+z+1$. Free-form input is easier to type initially; subsequent editing is also easier, for an expression may be typed as many short lines.

Extra white space can be forced into the output by several characters of various sizes. A tilde " ~ " gives a space equal to the normal word spacing in text; a circumflex gives half this much, and a tab charcter spaces to the next tab stop.

Spaces (or tildes, etc.) also serve to delimit pieces of the input. For example, to get

$$f(t)=2\pi \int \sin(\omega t)dt$$

we write

    f(t) = 2 pi int sin (omega t)dt

Here spaces are *necessary* in the input to indicate that *sin, pi, int,* and *omega* are special, and potentially worth special treatment. EQN looks up each such string of characters in a table, and if appropriate gives it a translation. In this case, *pi* and *omega* become their greek equivalents, *int* becomes the integral sign (which must be moved down and enlarged so it looks "right"), and *sin* is made roman, following conventional mathematical practice. Parentheses, digits and operators are automatically made roman wherever found.

Fractions are specified with the keyword *over:*

    a+b over c+d+e = 1

produces

$$\frac{a+b}{c+d+e}=1$$

Similarly, subscripts and superscripts are introduced by the keywords *sub* and *sup:*

$$x^2+y^2=z^2$$

is produced by

    x sup 2 + y sup 2 = z sup 2

The spaces after the 2's are necessary to mark the end of the superscripts; similarly the keyword *sup* has to be marked off by spaces or some equivalent delimiter. The return to the proper baseline is automatic. Multiple levels of subscripts or superscripts are of course allowed: "x sup y sup z" is x^{y^z}. The construct "something *sub* something *sup* something" is recognized as a special case, so "x sub i sup 2" is x_i^2 instead of $x_i{}^2$.

More complicated expressions can now be formed with these primitives:

$$\frac{\partial^2 f}{\partial x^2}=\frac{x^2}{a^2}+\frac{y^2}{b^2}$$

is produced by

    {partial sup 2 f} over {partial x sup 2} =
    x sup 2 over a sup 2 + y sup 2 over b sup 2

Braces {} are used to group objects together; in this case they indicate unambiguously what goes over what on the left-hand side of the expression. The language defines the precedence of *sup* to be higher than that of *over,* so no braces are needed to get the correct association on the right side. Braces can always be used when in doubt about precedence.

The braces convention is an example of

the power of using a recursive grammar to define the language. It is part of the language that if a construct can appear in some context, then *any expression* in braces can also occur in that context.

There is a *sqrt* operator for making square roots of the appropriate size: "sqrt a+b" produces $\sqrt{a+b}$, and

x = {−b +− sqrt{b sup 2 −4ac}} over 2a

is

$$x = \frac{-b \pm \sqrt{b^2-4ac}}{2a}$$

Since large radicals look poor on our typesetter, *sqrt* is not useful for tall expressions.

Limits on summations, integrals and similar constructions are specified with the keywords *from* and *to*. To get

$$\sum_{i=0}^{\infty} x_i \rightarrow 0$$

we need only type

sum from i=0 to inf x sub i −> 0

Centering and making the Σ big enough and the limits smaller are all automatic. The *from* and *to* parts are both optional, and the central part (e.g., the Σ) can in fact be anything:

lim from {x −> pi /2} (tan~x) = inf

is

$$\lim_{x \to \pi/2} (\tan x) = \infty$$

Again, the braces indicate just what goes into the *from* part.

There is a facility for making braces, brackets, parentheses, and vertical bars of the right height, using the keywords *left* and *right:*

left [x+y over 2a right]~=~1

makes

$$\left[\frac{x+y}{2a}\right] = 1$$

A *left* need not have a corresponding *right*, as we shall see in the next example. Any characters may follow *left* and *right*, but generally only various parentheses and bars are meaningful.

Big brackets, etc., are often used with another facility, called *piles*, which make vertical piles of objects. For example, to get

$$sign\,(x) \equiv \begin{cases} 1 & if\ x>0 \\ 0 & if\ x=0 \\ -1 & if\ x<0 \end{cases}$$

we can type

sign (x) ~=~ left {
  rpile {1 above 0 above −1}
  ~~lpile {if above if above if}
  ~~lpile {x>0 above x=0 above x<0}

The construction "left {" makes a left brace big enough to enclose the "rpile {...}", which is a right-justified pile of "above ... above ...". "lpile" makes a left-justified pile. There are also centered piles. Because of the recursive language definition, a pile can contain any number of elements; any element of a pile can of course contain piles.

Although EQN makes a valiant attempt to use the right sizes and fonts, there are times when the default assumptions are simply not what is wanted. For instance the italic *sign* in the previous example would conventionally be in roman. Slides and transparencies often require larger characters than normal text. Thus we also provide size and font changing commands: "size 12 bold {A~x~=~y}" will produce **A x = y**. *Size* is followed by a number representing a character size in points. (One point is 1/72 inch; this paper is set in 9 point type.)

If necessary, an input string can be quoted in "...", which turns off grammatical significance, and any font or spacing changes that might otherwise be done on it. Thus we can say

lim~ roman "sup" ~x sub n = 0

to ensure that the supremum doesn't become a superscript:

$$\lim \sup x_n = 0$$

Diacritical marks, long a problem in traditional typesetting, are straightforward:

$$\dot{x}+\hat{x}+\tilde{y}+\hat{X}+\ddot{Y}=\overline{z+Z}$$

is made by typing

x dot under + x hat + y tilde
+ X hat + Y dotdot = z+Z bar

There are also facilities for globally changing default sizes and fonts, for example for making viewgraphs or for setting chemical equations. The language allows for matrices, and for lining up equations at the same horizontal position.

Finally, there is a definition facility, so a user can say

define name "..."

at any time in the document; henceforth, any occurrence of the token "name" in an expression will be expanded into whatever was inside the double quotes in its definition. This lets users tailor the language to their own

specifications, for it is quite possible to redefine keywords like *sup* or *over*. Section 6 shows an example of definitions.

The EQN preprocessor reads intermixed text and equations, and passes its output to TROFF. Since TROFF uses lines beginning with a period as control words (e.g., ".ce" means "center the next output line"), EQN uses the sequence ".EQ" to mark the beginning of an equation and ".EN" to mark the end. The ".EQ" and ".EN" are passed through to TROFF untouched, so they can also be used by a knowledgeable user to center equations, number them automatically, etc. By default, however, ".EQ" and ".EN" are simply ignored by TROFF, so by default equations are printed in-line.

".EQ" and ".EN" can be supplemented by TROFF commands as desired; for example, a centered display equation can be produced with the input:

```
.ce
.EQ
x sub i = y sub i ...
.EN
```

Since it is tedious to type ".EQ" and ".EN" around very short expressions (single letters, for instance), the user can also define two characters to serve as the left and right delimiters of expressions. These characters are recognized anywhere in subsequent text. For example if the left and right delimiters have both been set to "#", the input:

Let #x sub i#, #y# and #alpha# be positive

produces:

Let x_i, y and α be positive

Running a preprocessor is strikingly easy on UNIX. To typeset text stored in file "f", one issues the command:

```
eqn f | troff
```

The vertical bar connects the output of one process (EQN) to the input of another (TROFF).

5. Language Theory

The basic structure of the language is not a particularly original one. Equations are pictured as a set of "boxes," pieced together in various ways. For example, something with a subscript is just a box followed by another box moved downward and shrunk by an appropriate amount. A fraction is just a box centered above another box, at the right altitude, with a line of correct length drawn between them.

The grammar for the language is shown below. For purposes of exposition, we have collapsed some productions. In the original grammar, there are about 70 productions, but many of these are simple ones used only to guarantee that some keyword is recognized early enough in the parsing process. Symbols in capital letters are terminal symbols; lower case symbols are non-terminals, i.e., syntactic categories. The vertical bar | indicates an alternative; the brackets [ ] indicate optional material. A TEXT is a string of non-blank characters or any string inside double quotes; the other terminal symbols represent literal occurrences of the corresponding keyword.

```
eqn  : box | eqn box

box  : text
     | { eqn }
     | box OVER box
     | SQRT box
     | box SUB box | box SUP box
     | [ L | C | R ]PILE { list }
     | LEFT text eqn [ RIGHT text ]
     | box [ FROM box ] [ TO box ]
     | SIZE text box
     | [ROMAN | BOLD | ITALIC] box
     | box [HAT | BAR | DOT | DOTDOT | TILDE]
     | DEFINE text text

list : eqn | list ABOVE eqn

text : TEXT
```

The grammar makes it obvious why there are few exceptions. For example, the observation that something can be replaced by a more complicated something in braces is implicit in the productions:

```
eqn : box | eqn box
box : text | { eqn }
```

Anywhere a single character could be used, *any* legal construction can be used.

Clearly, our grammar is highly ambiguous. What, for instance, do we do with the input

a over b over c ?

Is it

{a over b} over c

or is it

a over {b over c} ?

To answer questions like this, the grammar is supplemented with a small set of rules that describe the precedence and associativity of operators. In particular, we specify (more or less arbitrarily) that *over* associates to the left, so the first alternative above is the one chosen. On the other hand, *sub* and *sup* bind to the right,

because this is closer to standard mathematical practice. That is, we assume x^{a^b} is $x^{(a^b)}$, not $(x^a)^b$.

The precedence rules resolve the ambiguity in a construction like

a sup 2 over b

We define *sup* to have a higher precedence than *over*, so this construction is parsed as $\dfrac{a^2}{b}$ instead of $a^{\frac{2}{b}}$.

Naturally, a user can always force a particular parsing by placing braces around expressions.

The ambiguous grammar approach seems to be quite useful. The grammar we use is small enough to be easily understood, for it contains none of the productions that would be normally used for resolving ambiguity. Instead the supplemental information about precedence and associativity (also small enough to be understood) provides the compiler-compiler with the information it needs to make a fast, deterministic parser for the specific language we want. When the language is supplemented by the disambiguating rules, it is in fact LR(1) and thus easy to parse[5].

The output code is generated as the input is scanned. Any time a production of the grammar is recognized, (potentially) some TROFF commands are output. For example, when the lexical analyzer reports that it has found a TEXT (i.e., a string of contiguous characters), we have recognized the production:

    text   : TEXT

The translation of this is simple. We generate a local name for the string, then hand the name and the string to TROFF, and let TROFF perform the storage management. All we save is the name of the string, its height, and its baseline.

As another example, the translation associated with the production

    box   : box OVER box

is:

Width of output box =
  slightly more than largest input width
Height of output box =
  slightly more than sum of input heights
Base of output box =
  slightly more than height of bottom input box
String describing output box =
  move down;
  move right enough to center bottom box;
  draw bottom box (i.e., copy string for bottom box);
  move up; move left enough to center top box;
  draw top box (i.e., copy string for top box);
  move down and left; draw line full width;
  return to proper base line.

Most of the other productions have equally simple semantic actions. Picturing the output as a set of properly placed boxes makes the right sequence of positioning commands quite obvious. The main difficulty is in finding the right numbers to use for esthetically pleasing positioning.

With a grammar, it is usually clear how to extend the language. For instance, one of our users suggested a TENSOR operator, to make constructions like

$$_{m}^{l}T_{n\,i}^{k\,j}$$

Grammatically, this is easy: it is sufficient to add a production like

    box   : TENSOR { list }

Semantically, we need only juggle the boxes to the right places.

6. Experience

There are really three aspects of interest—how well EQN sets mathematics, how well it satisfies its goal of being "easy to use," and how easy it was to build.

The first question is easily addressed. This entire paper has been set by the program. Readers can judge for themselves whether it is good enough for their purposes. One of our users commented that although the output is not as good as the best hand-set material, it is still better than average, and much better than the worst. In any case, who cares? Printed books cannot compete with the birds and flowers of illuminated manuscripts on esthetic grounds, either, but they have some clear economic advantages.

Some of the deficiencies in the output could be cleaned up with more work on our part. For example, we sometimes leave too much space between a roman letter and an italic one. If we were willing to keep track of the fonts involved, we could do this better more of the

time.

Some other weaknesses are inherent in our output device. It is hard, for instance, to draw a line of an arbitrary length without getting a perceptible overstrike at one end.

As to ease of use, at the time of writing, the system has been used by two distinct groups. One user population consists of mathematicians, chemists, physicists, and computer scientists. Their typical reaction has been something like:

(1) It's easy to write, although I make the following mistakes...

(2) How do I do...?

(3) It botches the following things.... Why don't you fix them?

(4) You really need the following features...

The learning time is short. A few minutes gives the general flavor, and typing a page or two of a paper generally uncovers most of the misconceptions about how it works.

The second user group is much larger, the secretaries and mathematical typists who were the original target of the system. They tend to be enthusiastic converts. They find the language easy to learn (most are largely self-taught), and have little trouble producing the output they want. They are of course less critical of the esthetics of their output than users trained in mathematics. After a transition period, most find using a computer more interesting than a regular typewriter.

The main difficulty that users have seems to be remembering that a blank is a delimiter; even experienced users use blanks where they shouldn't and omit them when they are needed. A common instance is typing

    f(x sub i)

which produces

$$f(x_i)$$

instead of

$$f(x_i)$$

Since the EQN language knows no mathematics, it cannot deduce that the right parenthesis is not part of the subscript.

The language is somewhat prolix, but this doesn't seem excessive considering how much is being done, and it is certainly more compact than the corresponding TROFF commands. For example, here is the source for the continued fraction expression in Section 1 of this paper:

    a sub 0 + b sub 1 over
    {a sub 1 + b sub 2 over
     {a sub 2 + b sub 3 over
      {a sub 3 + ... }}}

This is the input for the large integral of Section 1; notice the use of definitions:

    define emx "{e sup mx}"
    define mab "{m sqrt ab}"
    define sa "{sqrt a}"
    define sb "{sqrt b}"
    int dx over {a emx - be sup -mx} ~=~
    left { lpile {
        1 over {2 mab} ~log~
            {sa emx - sb} over {sa emx + sb}
      above
        1 over mab ~ tanh sup -1 (sa over sb emx)
      above
        -1 over mab ~ coth sup -1 (sa over sb emx)
    }

As to ease of construction, we have already mentioned that there are really only a few person-months invested. Much of this time has gone into two things—fine-tuning (what is the most esthetically pleasing space to use between the numerator and denominator of a fraction?), and changing things found deficient by our users (shouldn't a tilde be a delimiter?).

The program consists of a number of small, essentially unconnected modules for code generation, a simple lexical analyzer, a canned parser which we did not have to write, and some miscellany associated with input files and the macro facility. The program is now about 1600 lines of C [6], a high-level language reminiscent of BCPL. About 20 percent of these lines are "print" statements, generating the output code.

The semantic routines that generate the actual TROFF commands can be changed to accommodate other formatting languages and devices. For example, in less than 24 hours, one of us changed the entire semantic package to drive NROFF, a variant of TROFF, for typesetting mathematics on teletypewriter devices capable of reverse line motions. Since many potential users do not have access to a typesetter, but still have to type mathematics, this provides a way to get a typed version of the final output which is close enough for debugging purposes, and sometimes even for ultimate use.

7. Conclusions

We think we have shown that it is possible to do acceptably good typesetting of mathematics on a phototypesetter, with an input language that is easy to learn and use and that satisfies many users' demands. Such a package can be implemented in short order, given a compiler-compiler

and a decent typesetting program underneath.

Defining a language, and building a compiler for it with a compiler-compiler seems like the only sensible way to do business. Our experience with the use of a grammar and a compiler-compiler has been uniformly favorable. If we had written everything into code directly, we would have been locked into our original design. Furthermore, we would have never been sure where the exceptions and special cases were. But because we have a grammar, we can change our minds readily and still be reasonably sure that if a construction works in one place it will work everywhere.

Acknowledgements

We are deeply indebted to J. F. Ossanna, the author of TROFF, for his willingness to modify TROFF to make our task easier and for his continuous assistance during the development of our program. We are also grateful to A. V. Aho for help with language theory, to S. C. Johnson for aid with the compiler-compiler, and to our early users A. V. Aho, S. I. Feldman, S. C. Johnson, R. W. Hamming, and M. D. McIlroy for their constructive criticisms.

References

[1] *A Manual of Style*, 12th Edition. University of Chicago Press, 1969. p 295.

[2] *Model C/A/T Phototypesetter*. Graphic Systems, Inc., Hudson, N. H.

[3] Ritchie, D. M., and Thompson, K. L., "The UNIX time-sharing system." *Comm. ACM 17*, 7 (July 1974), 365-375.

[4] Ossanna, J. F., TROFF User's Manual. Bell Laboratories Computing Science Technical Report 54, 1977.

[5] Aho, A. V., and Johnson, S. C., "LR Parsing." *Comp. Surv. 6*, 2 (June 1974), 99-124.

[6] B. W. Kernighan and D. M. Ritchie, *The C Programming Language*. Prentice-Hall, Inc., 1978.

UNIX Programming — Second Edition

Brian W. Kernighan

Dennis M. Ritchie

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

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

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

There is also an appendix which describes the standard I/O library in detail.

November 12, 1978

---

†UNIX is a Trademark of Bell Laboratories.

UNIX Programming — Second Edition

Brian W. Kernighan

Dennis M. Ritchie

Bell Laboratories
Murray Hill, New Jersey 07974

1. INTRODUCTION

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

The document collects material which is scattered throughout several sections of *The UNIX Programmer's Manual* [1] for Version 7 UNIX. There is no attempt to be complete; only generally useful material is dealt with. It is assumed that you will be programming in C, so you must be able to read the language roughly up to the level of *The C Programming Language* [2]. Some of the material in sections 2 through 4 is based on topics covered more carefully there. You should also be familiar with UNIX itself at least to the level of *UNIX for Beginners* [3].

2. BASICS

2.1. Program Arguments

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

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

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

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

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

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

2.2. 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. 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 need know nothing about where its input is coming from. This is also true if the input comes from another program via the pipe mechanism:

```
otherprog | prog
```

provides the standard input for `prog` from the standard output of `otherprog`.

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

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

```
prog >outfile
```

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

```
prog | otherprog
```

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

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

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

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

```
#include <stdio.h>

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

The line

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

should appear at the beginning of each source file. It causes the C compiler to read a file (*/usr/include/stdio.h*) of standard routines and symbols that includes the definition of `EOF`.

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

```
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 6 discusses status returns in more detail.

3. THE STANDARD I/O LIBRARY

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

In this section, we will discuss the basics of the standard I/O library. The appendix contains a more complete description of its capabilities.

3.1. File Access

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

        wc x.c y.c

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

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

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

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

        FILE *fp, *fopen();

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

The actual call to fopen in a program is

        fp = fopen(name, mode);

The first argument of fopen is the name of the file, as a character string. The second argument is the mode, also as a character string, which indicates how you intend to use the file. The only allowable modes are read ("r"), write ("w"), or append ("a").

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

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

        c = getc(fp)

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

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

puts the character c on the file fp and returns c. getc and putc return EOF on error.

When a program is started, three files are opened automatically, and file pointers are provided for them. These files are the standard input, the standard output, and the standard error output; the corresponding file pointers are called stdin, stdout, and stderr. Normally these are all connected to the terminal, but may be redirected to files or pipes as described in Section 2.2. stdin, stdout and stderr are pre-defined in the I/O library as the standard input, output and error files; they may be used anywhere an object of type FILE * can be. They are constants, however, *not* variables, so don't try to assign to them.

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

```c
#include <stdio.h>

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

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

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

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

3.2. Error Handling — Stderr and Exit

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

The program actually signals errors in another way, using the function exit to terminate program execution. The argument of exit is available to whatever process called it (see Section 6), so the success or failure of the program can be tested by another program that uses this one as a sub-process. By convention, a return value of 0 signals that all is well; non-zero values signal abnormal situations.

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

3.3. Miscellaneous I/O Functions

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

Normally output with putc, etc., is buffered (except to stderr); to force it out immediately, use fflush(fp).

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

The functions sscanf and sprintf are identical to fscanf and fprintf, except that the first argument names a character string instead of a file pointer. The conversion is done from the string for sscanf and into it for sprintf.

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

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

4. LOW-LEVEL I/O

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

4.1. File Descriptors

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

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

The file pointers discussed in section 3 are similar in spirit to file descriptors, but file descriptors are more fundamental. A file pointer is a pointer to a structure that contains, among other things, the file descriptor for the file in question.

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

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

```
prog <infile >outfile
```

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

4.2. Read and Write

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

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

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

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

The number of bytes to be read or written is quite arbitrary. The two most common values are 1, which means one character at a time ("unbuffered"), and 512, which corresponds to a physical blocksize on many peripheral devices. This latter size will be most efficient, but even character at a time I/O is not inordinately expensive.

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

```
#define    BUFSIZE    512 /* best size for PDP-11 UNIX */

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

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

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

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

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

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

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

c *must* be declared char, because read accepts a character pointer. The character being returned must be masked with 0377 to ensure that it is positive; otherwise sign extension may make it negative. (The constant 0377 is appropriate for the PDP-11 but not necessarily for other machines.)

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

```
#define    CMASK      0377 /* for making char's > 0 */
#define    BUFSIZE    512

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

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

4.3. Open, Creat, Close, Unlink

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

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

```
int fd;

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

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

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

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

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

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

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

```
#define NULL 0
#define BUFSIZE 512
#define PMODE 0644 /* RW for owner, R for group, others */

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

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

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

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

As we said earlier, there is a limit (typically 15-25) on the number of files which a program may have open simultaneously. Accordingly, any program which intends to process many files must be prepared to re-use file descriptors. The routine close breaks the connection between a file descriptor and an open file, and frees the file descriptor for use with some other file. Termination of a program via exit or return from the main program closes all open files.

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

4.4. Random Access — Seek and Lseek

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

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

forces the current position in the file whose descriptor is fd to move to position offset, which is taken relative to the location specified by origin. Subsequent reading or writing will begin at that position. offset is a long; fd and origin are int's. origin can be 0, 1, or 2 to specify that offset is to be measured from the beginning, from the current position, or from the end of the file respectively. For example, to append to a file, seek to the end before writing:

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

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

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

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

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

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

In pre-version 7 UNIX, the basic entry point to the I/O system is called seek. seek is identical to lseek, except that its offset argument is an int rather than a long. Accordingly, since PDP-11 integers have only 16 bits, the offset specified for seek is limited to 65,535; for this reason, origin 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 origin as if it were 0, 1, or 2 respectively. Thus to get to an arbitrary place in a large file requires two seeks, first one which selects the block, then one which has origin equal to 1 and moves to the desired byte within the block.

4.5. 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.

5. PROCESSES

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

5.1. The "System" Function

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

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

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

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

5.2. Low-Level Process Creation — Execl and Execv

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

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

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

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

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

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

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

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

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

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

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

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

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

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

5.3. Control of Processes — Fork and Wait

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

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

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

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

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

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

```
int status;

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

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

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

return meaningful status.

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

5.4. Pipes

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

```
ls | pr
```

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

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

```
int   fd[2];

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

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

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

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

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

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

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

```
#include <stdio.h>

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

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

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

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

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

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

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

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

```
#include <signal.h>

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

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

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

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

6. SIGNALS — INTERRUPTS AND ALL THAT

This section is concerned with how to deal gracefully with signals from the outside world (like interrupts), and with program faults. Since there's nothing very useful that can be done from within C about program faults, which arise mainly from illegal memory references or from execution of peculiar instructions, we'll discuss only the outside-world signals: *interrupt*, which is sent when the DEL character is typed; *quit*, generated by the FS character; *hangup*, caused by hanging up the phone; and *terminate*, generated by the *kill* command. When one of these events occurs, the signal is sent to *all* processes which were started from the corresponding terminal; unless other arrangements have been made, the signal terminates the process. In the *quit* case, a core image file is written for debugging purposes.

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

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

causes interrupts to be ignored, while

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

restores the default action of process termination. In all cases, `signal` returns the previous value of the signal. The second argument to `signal` may instead be the name of a function (which has to be declared explicitly if the compiler hasn't seen it already). In this case, the named routine will be called when the signal occurs. Most commonly this facility is used to

allow the program to clean up unfinished business before terminating, for example to delete a temporary file:

```
#include <signal.h>

main()
{
        int onintr();

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

        /* Process ... */   .

        exit(0);
}

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

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

The solution, shown above, is to test the state of interrupt handling, and to continue to ignore interrupts if they are already being ignored. The code as written depends on the fact that `signal` returns the previous state of a particular signal. If signals were already being ignored, the process should continue to ignore them; otherwise, they should be caught.

A more sophisticated program may wish to intercept an interrupt and interpret it as a request to stop what it is doing and return to its own command-processing loop. Think of a text editor: interrupting a long 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:

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

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

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

        /* main processing loop */
}
```

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

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

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

There is one difficulty associated with this approach. Suppose the program is reading the terminal when the interrupt is sent. The specified routine is duly called; it sets its flag and returns. If it were really true, as we said above, that "execution resumes at the exact point it was interrupted," the program would continue reading the terminal until the user typed another line. This behavior might well be confusing, since the user might not know that the program is reading; he presumably would prefer to have the signal take effect instantly. The method chosen to resolve this difficulty is to terminate the terminal read when execution resumes after the signal, returning an error code which indicates what happened.

Thus programs which catch and resume execution after signals should be prepared for "errors" which are caused by interrupted system calls. (The ones to watch out for are reads from a terminal, `wait`, and `pause`.) A program whose `onintr` 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() == EOF)
    if (intflag)
        /* EOF caused by interrupt */
    else
        /* true end-of-file */
```

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

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

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

```
#include <signal.h>

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

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

As an aside on declarations, the function `signal` obviously has a rather strange second argument. It is in fact a pointer to a function delivering an integer, and this is also the type of the signal routine itself. The two values `SIG_IGN` and `SIG_DFL` have the right type, but are chosen so they coincide with no possible actual functions. For the enthusiast, here is how they are defined for the PDP-11; the definitions should be sufficiently ugly and nonportable to encourage use of the include file.

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

References

[1] K. L. Thompson and D. M. Ritchie, *The UNIX Programmer's Manual*, Bell Laboratories, 1978.

[2] B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Prentice-Hall, Inc.. 1978.

[3] B. W. Kernighan, "UNIX for Beginners — Second Edition." Bell Laboratories, 1978.

Appendix — The Standard I/O Library

D. M. Ritchie

Bell Laboratories
Murray Hill, New Jersey 07974

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

1. It must be as efficient as possible, both in time and in space, so that there will be no hesitation in using it no matter how critical the application.

2. It must be simple to use, and also free of the magic numbers and mysterious calls whose use mars the understandability and portability of many programs using older packages.

3. The interface provided should be applicable on all machines, whether or not the programs which implement it are directly portable to other systems, or to machines other than the PDP-11 running a version of UNIX.

1. General Usage

Each program using the library must have the line

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

which defines certain macros and variables. The routines are in the normal C library, so no special library argument is needed for loading. All names in the include file intended only for internal use begin with an underscore _ to reduce the possibility of collision with a user name. The names intended to be visible outside the package are

stdin   The name of the standard input 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 I/O buffer supplied by the user. See setbuf, below.

getc, getchar, putc, putchar, feof, ferror, fileno
    are defined as macros. Their actions are described below; they are mentioned here to point out that it is not possible to redeclare them and that they are not actually functions; thus, for example, they may not have breakpoints set on them.

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

2. Calls

FILE *fopen(filename, type) char *filename, *type;
    opens the file and, if needed, allocates a buffer for it. filename is a character string specifying the name. type is a character string (not a single character). It may be "r", "w", or "a" to indicate intent to read, write, or append. 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, so it can be pointed to, passed as an argument, etc.

`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;`

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 I/O 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; however, `stderr` always starts off unbuffered and remains so unless `setbuf` is used, or unless it is reopened.

`exit(errcode);`

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

`feof(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 *fgets(s, n, ioptr) char *s; FILE *ioptr;`

reads up to n-1 characters from the stream `ioptr` into the character pointer `s`. The read terminates with a newline character. The newline character is placed in the buffer followed by a null character. `fgets` returns the first argument, or `NULL` if error or end-of-file occurred.

`fputs(s, ioptr) char *s; FILE *ioptr;`

writes the null-terminated string (character array) `s` on the stream `ioptr`. No newline 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(3) of the *UNIX Programmer's Manual.*

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

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

```
fread(ptr, sizeof(*ptr), nitems, ioptr) FILE *ioptr;
```
reads nitems of data beginning at ptr from file ioptr. No advance notification that binary I/O 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. A "word" is 16 bits on the PDP-11.

```
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 I/O has started. If buf is NULL, the stream will be unbuffered. Otherwise the buffer supplied will be used. It must be 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 *malloc(num);`
allocates `num` bytes. The pointer returned is sufficiently well aligned to be usable for any purpose. `NULL` is returned if no space is available.

`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 whose definitions may be obtained by including `<ctype.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, newline, carriage return, vertical tab, form feed, space.

`ispunct(c)` returns non-zero if the argument is any punctuation character, i.e., not a space, letter, digit or control character.

`isalnum(c)` returns non-zero if the argument is a letter or a digit.

`isprint(c)` returns non-zero if the argument is printable — a letter, digit, or punctuation character.

`iscntrl(c)` returns non-zero if the argument is a control character.

`isascii(c)` returns non-zero if the argument is an ascii character, i.e., less than octal 0200.

`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`.

The C Environment of UNIX/TS

Andrew R. Koenig

Bell Laboratories
Murray Hill, New Jersey 07974

1. INTRODUCTION

This document describes differences users may encounter when changing to UNIX†/TS from the various so-called "UNIX Sixth Edition" C compilers. The document is intended as a conversion aid, so the emphasis is on incompatibilities, rather than new facilities.

Note that this document is only a guide; refer to the *UNIX/TS User's Manual* for complete information. This version of this document supersedes all previous versions thereof.

2. LIBRARY CHANGES

The changes that are most likely to be noticed are in the run-time library. The "Standard I/O Library" has been incorporated into */lib/libc.a* along with the contents of */lib/liba.a*; this latter library is gone. *Printf* has been rewritten into portable C; there are a few incompatibilities with the old version. Finally, there are a number of smaller changes and incompatibilities.

2.1 Environments

The system now makes available to the user program a table of *environment variables*. Each variable has a name and a value; both name and value are character strings. The values of environment variables are preserved across *fork* and *exec*; they can also be altered easily using the shell and somewhat less easily using the new *execle* and *execve* system calls.

The new *getenv* function can be used to retrieve the value of an environment variable.

2.2 The Standard I/O Library

In the past, there were two I/O libraries available. One was documented by *A New Input-Output Package* (Ritchie), and was made available through the **-lS** loader option. The other, older one was made available whenever a C program was being compiled; it was characterized, among other things, by use of the names *fin* and *fout* to control disposition of standard input and output files.

The older library has now vanished, along with the **-lS** option. All programs will receive the new I/O library without any explicit action. In addition, the libraries obtained by **-lc** and **-la** have been merged; this combined library is accessed (where needed) by **-lc** or **-l**.

2.3 Printf

In the interests of portability, *printf* has been rewritten into portable C. This results in load modules some 1800 bytes larger than previous versions.

The correct way to write a long integer is now **%ld** or **%lo**; the previous forms **D** and **O** are going away. The purpose of this is to permit **X**, **E**, and **G** format codes for indicating that the letters produced by the format code are to appear in upper case.

The **%r** format code has been removed; if you don't know what it was, you don't want to know.

2.4 Scanf

---

† UNIX is a Trademark of Bell Laboratories.

2.4.1 White space

The way in which *scanf* treats white space has changed slightly. No longer is it the case that *scanf* will skip white space in the input for each character in the format. Rather, a space, tab, or new-line in the format will match optional white space in the input. Thus:

> "alpha = %d"

will match any of

> alpha=12
> alpha =12
> alpha= 12
> alpha = 12

but not

> a lpha=12

as was formerly the case. Note that this change may require white space to be inserted in format strings of formerly working programs to maintain compatibility.

2.4.2 Character class formats

A character class format item (such as "%[0123456789]") is now permitted to match a null string. Thus,

> scanf (":%[^:]:", x);

will no longer fail when presented with

> ::

2.5 Mathematical Routines

The mathematical subroutines have been moved to a separate library obtainable by the -lm option. Declarations for these routines can be obtained in <**math.h**>.

2.6 Character class routines

The routines that test character class (*isdigit*, etc.) are no longer defined in <**stdio.h**>; rather, they are defined in <**ctype.h**>. Thus, a line of the form

> #include <ctype.h>

will have to be added to those programs which use *isdigit*, *isupper*, *islower*, and their relatives. The domain of the character class routines has been extended to match the range of *getc*: -1 through 255. The character class routine *isprint* has been revised to conform to its documentation; a space is now considered a printable character. To determine if a character has a graphic representation, use the (new) function *isgraph*.

2.7 Character Conversion Routines

The routines *toupper* and *tolower* have had their domain extended to the range of *getc*: *toupper* will return its argument unchanged if that argument is not a lower-case letter, and *toupper* will return its argument if it is not an upper-case letter. This change required rewriting *toupper* and *tolower* as true subroutines, rather than macros; for applications where efficiency is paramount and the argument is already known to be a letter of the appropriate case, the original macros have been renamed _*toupper* and _*tolower*.

2.8 Error Recovery

The library now incorporates two new routines, *ssignal* and *gsignal*. In the future, these routines will be used by other routines in the library to cause automatic program termination on detection of various common errors, with the possibility of finer control as a user option.

This description is deliberately vague, as the facility is still in the planning stage.

2.9 Time of Day

There is a new function *tzset*. It is called with no arguments, and looks for an environment variable **TZ**. This variable is expected to be in the form **EST***n* or **EST***n***EDT**, where *n* is a string of digits with an optional negative sign and represents the difference between the local time zone and GMT, surrounded by the names of the local and (optional) daylight time zones. If *tzset* finds an environment variable **TZ** in this form, it sets the time zone parameters *timezone*, *tzname*, and *daylight* appropriately. *Tzset* is now called automatically by *asctime*, so it usually need not be called by the user.

Note also that the variable *timezone* is now a **long**, so programs referencing it will have to be changed slightly.

2.10 Miscellaneous

2.10.1 chown

Chown now takes three arguments: the file name, the new owner, and the new group. This is necessary because owner and group can now each be up to 16 bits.

2.10.2 tell

Tell is gone; *lseek* instead returns a value indicating the location sought.

2.10.3 setexit and reset

Setexit and *reset* are gone; their function is taken over by *setjmp* and *longjmp*. These new routines provide all the facilities of *setexit* and *reset* in a more general form.

2.10.4 nargs

Nargs is gone. There is no replacement routine, as *nargs* cannot be made to work with separate I and D space.

2.10.5 String routines

Strcatn, *strcpyn*, *strcmpn*, *index*, and *rindex* have been renamed *strncat*, *strncpy*, *strncmp*, *strchr*, and *strrchr*, respectively. This follows the recommendations of the C Standards Task Force, and also allows compatibility with systems that require distinct external names to differ within their first six characters.

2.10.6 Effective user and group ID

There are two new routines, *geteuid* and *getegid*, which return the effective user and group ID, rather than the real user and group ID.

2.10.7 time

The *time* routine now returns a **long** value; it will also store a copy of the value in the (long) location addressed by its argument unless that argument is **(long *)0**.

2.10.8 The password file

The format of */etc/passwd* has changed slightly with the introduction of UNIX/TS; this change is reflected in the various routines which extract information from */etc/passwd*. In addition, a new file, */etc/group*, has been created to hold information about group access privileges. This file is searched by a new set of routines.

The names of the routines under discussion are:

```
endpwent        endgrent
getpwent        getgrent
getpwnam
getpwuid        getgrgid
setpwent        setgrent
```

3. THE LANGUAGE

3.1 The Preprocessor

John Reiser has rewritten the C preprocessor. The new one is largely compatible with the old one, and much faster, but there are a few changes.

3.1.1 General

Symbols defined on the command line by **-D**foo are defined as **1**, i.e., as if they had been defined by

#define foo 1

or

-Dfoo=1

This means that names automatically defined by the preprocessor (specifically *unix* and *pdp11*) cannot be used as identifiers in the program without naming them in **#undef** statements or using the **-U** preprocessor option.

The directory search order for **#include** requests is:

1. the directory of the file which contains the **#include** request (e.g. **#include** is relative to the file being scanned when the request is made), for statements of the form

    #include "*name*"

2. the directories specified by **-I**, in left-to-right order (as usual, the null string can be used to name the current directory)
3. the standard directory(s) (which for the UNIX system is */usr/include*)

An unescaped new-line terminates a character constant or quoted string.

An escaped new-line (a backslash immediately followed by a new-line) may be used in the body of a **#define** statement to continue the definition onto the next line. The escaped new-line is not included in the macro body.

Comments are uniformly removed (except if the argument **-C** is specified). They are also ignored, except that a comment terminates a token. Thus

foo/* la di da */bar

may expand 'foo' and 'bar' but will never expand 'foobar'. If neither 'foo' nor 'bar' is a macro then the output is the string 'foobar', even if the preprocessor name 'foobar' is defined as something else. The file

#define foo(a,b)b/**/a
foo(1,2)

produces '21' because the comment causes a break which enables the recognition of 'b' and 'a' as formals in the string "b/**/a".

Macro formal parameters are recognized in **#define** bodies even inside character constants and quoted strings. The output from

```
#define foo(a) '\\a'
foo(bar)
```

is the seven characters " '\\bar'". Macro names are not recognized inside character constants or quoted strings during the regular scan. Thus

```
#define foo bar
printf("foo");
```

does not expand 'foo' in the second line, because it is inside a quoted string which is not part of a #**define** macro definition.

Macros are not expanded while processing a #**define** or #**undef**. Thus

```
#define foo bletch
#define bar foo
#undef foo
bar
```

produces 'foo'. The token appearing immediately after an #**ifdef** or #**ifndef** is not expanded (of course!).

Macros are not expanded during the scan which determines the actual parameters to another macro call. Thus

```
#define foo(a,b)b a
#define bar hi
foo(bar,
#define bar bye
)
```

produces " bye" (and warns about the redefinition of 'bar').

3.1.2 Bugs fixed

1. "1.e4" is recognized as a floating-point number, rather than as an opportunity to expand the possible macro name "e4".
2. Any kind and amount of white space (space, tab, line-feed, vertical tab, form-feed, carriage return) is allowed between a macro name and the left parenthesis which introduces its actual parameters.
3. The comma operator is legal in preprocessor #**if** statements.
4. Macros with parameters are legal in preprocessor #**if** statements.
5. Single-character character constants are legal in preprocessor #**if** statements.
6. Line-feeds are put out in the proper place when a multi-line comment is not passed through to the output.
7. The following example expands to "# # #" :

    ```
    #define foo #
    foo foo foo
    ```

8. If the -R flag is not specified then the invocation of some recursive macros is trapped and the recursion forcibly terminated with an error message. The recursions that are trapped are the ones in which the nesting level is non-decreasing from some point on. In particular,

    ```
    #define a a
    a
    ```

    will be detected. (Use "#undef a" if that is what you want.)
9. The recursion

```
#define a c b
#define b c a
#define c foo
a
```

will not be detected because the nesting level decreases after each expansion of "c".

10. The -R flag specifically allows recursive macros and recursion will be strictly obeyed (to the extent that space is available). Assuming that -R is specified:

```
#define a a
a
```

causes an infinite loop with very little output. The tail recursion

```
#define a <b
#define b >a
a
```

causes the string " < > " to be output infinitely many times. The non-tail recursion

```
#define a b>
#define b a<
a
```

complains "too much pushback", dumps the "pushback", and continues (again, infinitely).

3.1.3 Stylistic choice

1. Nothing (not even line-feeds) is output while a false #if, #ifdef, or #ifndef is in effect. Thus when all conditions become true a line of the form '# 12345 "foo.c"' is output.
2. Error and warning messages always appear on standard error (file descriptor 2).
3. Mismatch between the number of formals and actuals in a macro call produces only a warning, and not an error. Excess actuals are ignored; missing actuals are turned into null strings.

3.1.4 Incompatibility

The virgule '/' in "a=/*b" is interpreted as the first character of the pair "/*" which introduces a comment, rather than as the second character of the divide-and-replace operator "=/". This incompatibility reflects the recent change in the C language which made "a/=*b" the legal way to write such a statement if the meaning "a=a/ *b" is intended.

3.2 The Compiler

3.2.1 Enumerated Data Types

Enumerated data types are here, though not yet documented, so that **enum** is now a keyword.

3.2.2 Unsigned numbers

The value returned by *sizeof* is now *unsigned* rather than *int*, so care must be exercised in the use of **sizeof** in a few strange cases. For example, the following no longer works:

```
if (n < - sizeof (x)) { ... }
```

because unary - is meaningless when applied to an **unsigned** value.

3.2.3 Structure and Union Assignments

It is now possible to assign structures and pass them as arguments and results of procedures. This feature is not new in the latest release, but it is sufficiently important that it is worth noting anyway.

4. SOURCE STRUCTURE

The new preprocessor and changes in the library make the source structure of this new release of C different from previous versions.

4.1 The Compiler

The new preprocessor is comprised of three source modules: *cpp.c*, *cpy.y*, and *yylex.c*. *Cpy.y* should be processed by *yacc* to produce *cpy.c*; this and *cpp.c* should then be compiled together to produce the preprocessor. Despite its name, *yylex.c* does not involve using *lex*, and it is not directly compiled; rather, it is named by #**includes** in the other modules.

4.2 The Library

The source for the mathematical routines in the C library is now in */usr/src/lib/libm*. The source in */usr/src/lib/libc* is now organized in five subdirectories:

1. *crt*, which contains run-time routines that are invoked by generated object code without ever being explicitly referenced by the programmer. These routines are largely in assembler language, and do things like **long** multiplication and division.
2. *csu*, which contains routines that are explicitly referenced by the *cc* command; these routines are used for run-time initialization.
3. *gen*, which contains those routines described in section 3 of the manual that are not part of the "standard I/O package",
4. *stdio*, which contains those routines described in section 3 of the manual that *are* part of the "standard I/O package", and
5. *sys*, which contains the routines described in section 2 of the manual. These routines are all in assembler language, and are interfaces between the C language and the UNIX system calls.

The other files in the */usr/src/lib/libc* directory are used as part of the installation procedures. *Order.in* and *order.out* are used to define the ordering of the modules in */lib/libc.a*, and *libc.rc* is a command file to recompile the library.

May 1979
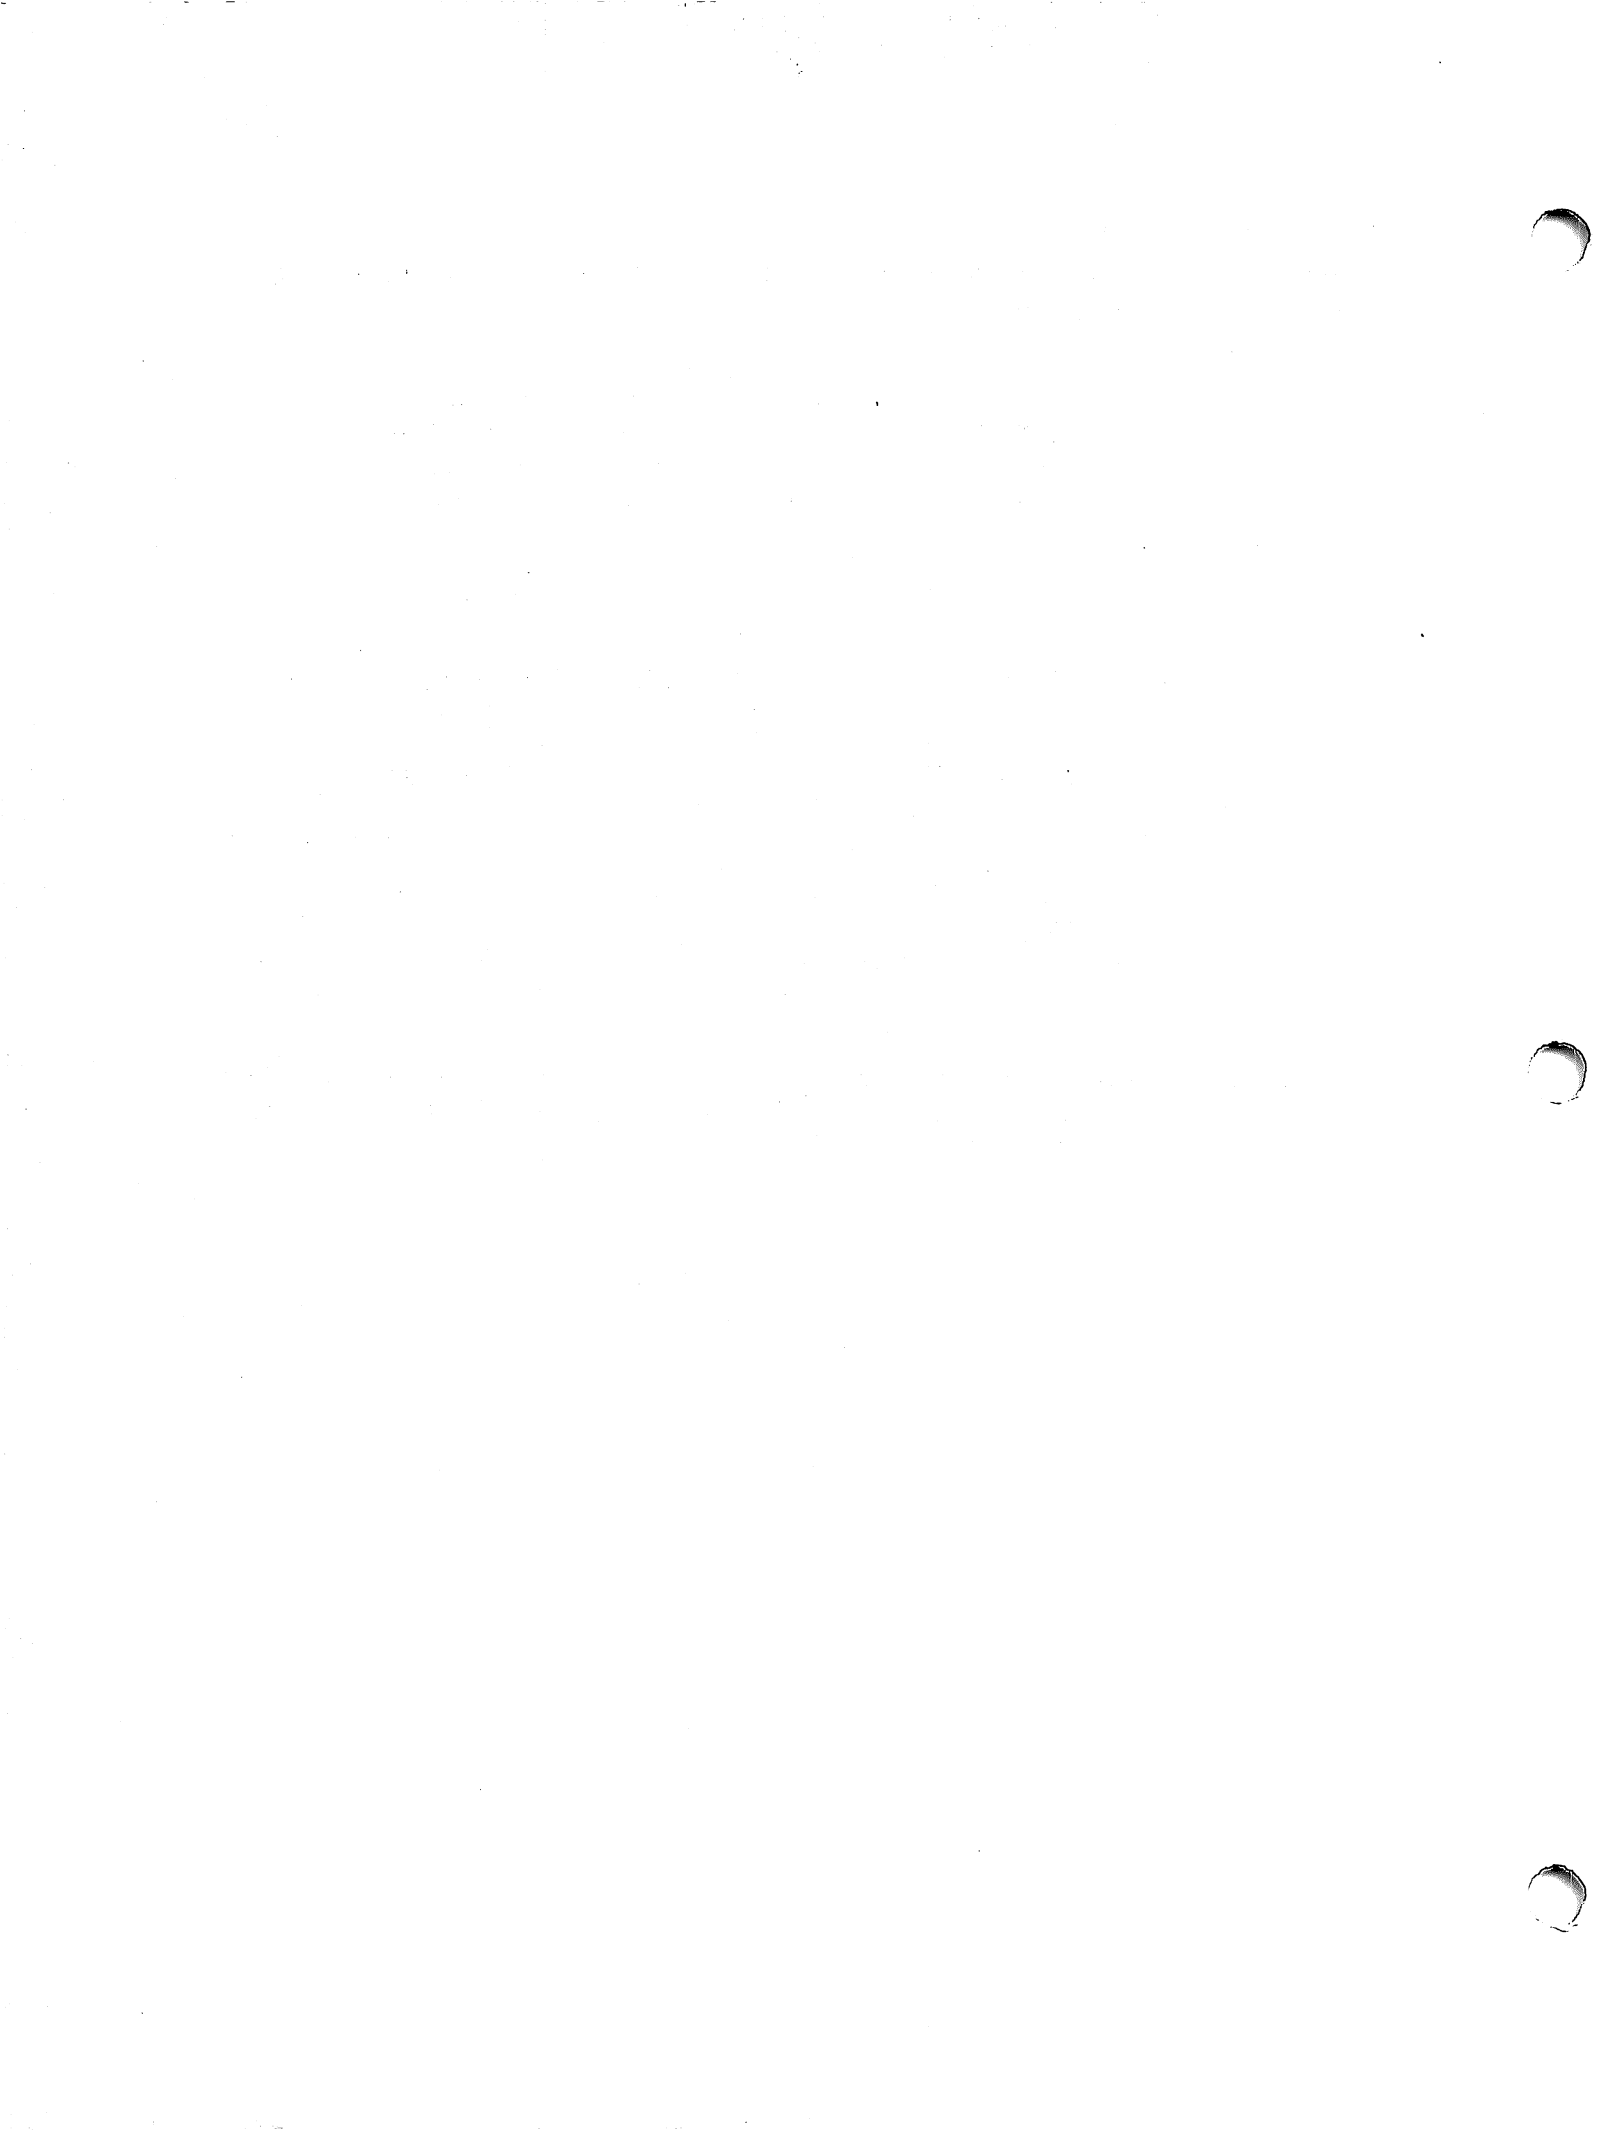
The C Programming Language — Reference Manual

Dennis M. Ritchie

Bell Laboratories, Murray Hill, New Jersey

This manual is reprinted, with minor changes, from *The C Programming Language*, by Brian W. Kernighan and Dennis M. Ritchie, Prentice-Hall, Inc., 1978.

1. Introduction

This manual describes the C language on the DEC PDP-11, the DEC VAX-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

There are six classes of tokens: identifiers, keywords, constants, strings, operators, and other separators. Blanks, tabs, newlines, and comments (collectively, "white space") as described below are ignored except as they serve to separate tokens. Some white 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 a letter. The underscore _ counts as a letter. Upper and lower case letters are different. No more than the first eight characters are significant, although more may be used. External identifiers, which are used by various assemblers and loaders, are more restricted:

DEC PDP-11	7 characters, 2 cases
DEC VAX-11	8 characters, 2 cases
Honeywell 6000	6 characters, 1 case
IBM 360/370	7 characters, 1 case
Interdata 8/32	8 characters, 2 cases

2.3 Keywords

The following identifiers are reserved for use as keywords, and may not be used otherwise:

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 words fortran and asm.

2.4 Constants

There are several kinds of constants, as listed below. Hardware characteristics which affect sizes are summarized in §2.6.

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 is taken to be long; an octal or hex constant which exceeds the largest unsigned machine integer 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. As discussed below, on some machines integer and long values may be considered identical.

2.4.3 Character constants

A character constant is a character enclosed in single quotes, as in 'x'. The value of a character constant is the numerical value of the character in the machine's character set.

Certain non-graphic characters, the single quote ' and the backslash \, may be represented according to the following table of escape sequences:

newline	NL (LF)	\n
horizontal tab	HT	\t
backspace	BS	\b
carriage return	CR	\r
form feed	FF	\f
backslash	\	\\
single quote	'	\'
bit pattern	*ddd*	*ddd*

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 is ignored.

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, as in "...". A string has type "array of characters" and storage class static (see §4 below) and is initialized with the given characters. All strings, even when written identically, are distinct. 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 double quote 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 newline are ignored.

2.6 Hardware characteristics

The following table summarizes certain hardware properties which vary from machine to machine. Although these affect program portability, in practice they are less of a problem than might be thought *a priori*.

	DEC PDP-11	Honeywell 6000	IBM 370	Interdata 8/32
	ASCII	ASCII	EBCDIC	ASCII
char	8 bits	9 bits	8 bits	8 bits
int	16	36	32	32
short	16	36	16	16
long	32	36	32	32
float	32	36	32	32
double	64	72	64	64
range	$\pm 10^{\pm 38}$	$\pm 10^{\pm 38}$	$\pm 10^{\pm 76}$	$\pm 10^{\pm 76}$

The VAX-11 is identical to the PDP-11 except that integers have 32 bits.

3. Syntax notation

In the syntax notation used in this manual, syntactic categories are indicated by *italic* type, and literal words and characters in **bold** type. Alternative categories are listed on separate lines. An optional terminal or non-terminal symbol is indicated by the subscript "opt," so that

$$\{ \ expression_{opt} \ \}$$

indicates an optional expression enclosed in braces. The syntax is summarized in §18.

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 (§9.2), 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 from that character set 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.

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.

Unsigned integers, declared unsigned, obey the laws of arithmetic modulo 2^n where n is the number of bits in the representation. (On the PDP-11, unsigned long quantities are not supported.)

Single-precision floating point (float) and double-precision floating point (double) may be synonymous in some implementations.

Because objects of the foregoing 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 shorter integer to a longer always involves sign extension; 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. Of the machines treated by this manual, only the PDP-11 sign-extends. On the PDP-11, character variables range in value from −128 to 127; the characters of the ASCII alphabet are all positive. A character constant specified with an octal escape suffers sign extension and may appear negative; for example, '\377' has the value −1.

When a longer integer is converted to a shorter or to a char, it is truncated on the left; excess bits are simply discarded.

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; in particular the direction of truncation of negative numbers varies from machine to machine. 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 is the least unsigned integer congruent to the signed integer (modulo $2^{wordsize}$). In a 2's complement representation, 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, for example, 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 grammar of §18.

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. The order in which side effects take place is unspecified. 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.

The handling of overflow and divide check in expression evaluation is machine-dependent. All existing implementations of C ignore integer overflows; treatment of division by 0, and all floating-point exceptions, varies between machines, and is usually adjustable by a library function.

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-list$_{opt}$*)
>   *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. If the type of the identifier is "array of ...", however, 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 constant is a primary expression. Its type may be int, long, or double depending on its form. Character constants have type int; floating constants are double.

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`; and as usual, array names are converted to pointers. No other conversions are performed automatically; in particular, the compiler does not compare the types of actual arguments with those of formal arguments. If conversion is needed, use a cast; see §7.2, 8.7.

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. An array name is a pointer expression. 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 a 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 2^n, where n is the number of bits in an `int`. There is no unary + operator.

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 ++x is equivalent to x+=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 prefix ++ 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. This construction is called a *cast*. Type names are 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 by the compiler.

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. On all machines covered by this manual, the remainder has the same sign as the dividend. It is always true that `(a/b)*b + a%b` is equal to `a` (if `b` is not 0).

The binary `%` operator yields the remainder from the division of the first expression by the second. The usual arithmetic conversions are performed. The operands must not be `float`.

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 for pointers.

The `+` operator is associative and expressions with several additions at the same level may be rearranged by the compiler.

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 greater than or equal to the length of the object in bits.

> *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 right 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; 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 bitwise 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 bitwise 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 involving | may be rearranged. The usual arithmetic conversions are performed; the result is the bitwise 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

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 compilers currently allow 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 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 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 other 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 improvements in code generation 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 never automatic.

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.

> > `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.

Now imagine a declaration

                    T D1

where T is a type-specifier (like int, etc.) and D1 is a declarator. Suppose this declaration makes the identifier have type "... T," where the "..." is empty if D1 is just a plain identifier (so that the type of x in "int x" is just int). Then if D1 has the form

                    *D

the type of the contained identifier is "... pointer to T."
    If D1 has the form

                    D()

then the contained identifier has the type "... function returning T."
    If D1 has the form

                    D[constant-expression]

or

                    D[]

then the contained identifier has type "... array of T." 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.) When several "array of" specifications are adjacent, a multi-dimensional array is created; the constant expressions which specify the bounds of the arrays may be missing only for the first member of the sequence. This elision is useful when the array is external and the actual definition, which allocates storage, is given elsewhere. The first 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, unions 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 or union 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; it returns an integer.
    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 3×5×7. 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.

```
struct-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; therefore, there may be 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. Fields are assigned right-to-left on the PDP-11, left-to-right on other machines.

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. In all implementations, there are no arrays of fields, and the address-of operator & may not be applied to them, so that there are no pointers to fields.

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 illegal to declare a structure or union which contains an instance of itself, but a structure or union may contain 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 declaration

```
struct tnode s, *sp;
```

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; and

```
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* }

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 constants, and previously declared variables and functions.

Static and external variables which are not initialized are guaranteed to start off as 0; automatic and register variables which are not initialized are guaranteed to start off as garbage.

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.

Braces may be elided as follows. If the initializer begins with a left brace, then the succeeding comma-separated list of initializers initializes 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 characters 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 by means of a cast, 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-expression$_{opt}$*]

To avoid ambiguity, in the construction

 (*abstract-declarator*)

the abstract-declarator is required to be non-empty. 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]
int *()
int (*)()
```

name respectively the types "integer," "pointer to integer," "array of 3 pointers to integers," "pointer to an array of 3 integers," "function returning pointer to integer," and "pointer to function returning an integer."

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.

> *typedef-name:*
>> *identifier*

Within the scope of a declaration involving typedef, each identifier appearing as part of any declarator therein become syntactically equivalent to the type keyword naming the type associated with the identifier in the way described in §8.4. 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 object.

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, after which 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, `extern` 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 `else`-less `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-1*$_{opt}$` ; `*expression-2*$_{opt}$` ; `*expression-3*$_{opt}$` ) `*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 often 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 labeled 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, but 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: ;            contin: ;          contin: ;
}                    } while (...);      }
```

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.

9.12 Labeled 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.

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 structures, unions and functions cannot be passed to a function, it is useless to declare a formal parameter to be a structure, union or function (pointers to such objects are of course permitted).

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 source 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. Members and tags follow the same scope rules as other identifiers. `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 refers to an identifier declared 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 size specified in the definition are compatible with those specified by each function which references the data.

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 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. Functions may be declared `static`.

---

†It is agreed that the ice is thin here.

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.

#include's 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 line number of the next source line 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. 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 already 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 must be declared explicitly in the calling routine since its appearance in g(f) 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 $i \times j \times \cdots \times k$, then E appearing in an expression is converted to a pointer to an $(n-1)$-dimensional array with rank $j \times \cdots \times k$. 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.

14.4 Explicit pointer conversions

Certain conversions involving pointers are permitted but have implementation-dependent aspects. They are all specified by means of an explicit type-conversion operator, §§7.2 and 8.7.

A pointer may be converted to any of the integral types large enough to hold it. Whether an int or long is required is machine dependent. The mapping function is also machine dependent, but is intended to be unsurprising to those who know the addressing structure of the machine. Details for some particular machines are given below.

An object of integral type may be explicitly converted to a pointer. The mapping always carries an integer converted from a pointer back to the same pointer, but is otherwise machine dependent.

A pointer to one type may be converted to a pointer to another type. The resulting pointer may cause addressing exceptions upon use if the subject pointer does not refer to an object suitably aligned in storage. It is guaranteed that a pointer to an object of a given size may be converted to a pointer to an object of a smaller size and back again without change.

For example, a storage-allocation routine might accept a size (in bytes) of an object to allocate, and return a char pointer; it might be used in this way.

```
extern char *alloc();
double *dp;

dp = (double *) alloc(sizeof(double));
*dp = 22.0 / 7.0;
```

alloc must ensure (in a machine-dependent way) that its return value is suitable for conversion to a pointer to double; then the *use* of the function is portable.

The pointer representation on the PDP-11 corresponds to a 16-bit integer and is measured in bytes. chars have no alignment requirements; everything else must have an even address.

On the Honeywell 6000, a pointer corresponds to a 36-bit integer; the word part is in the left 18 bits, and the two bits that select the character in a word just to their right. Thus char pointers are measured in units of 2^{16} bytes; everything else is measured in units of 2^{18} machine words. double quantities and aggregates containing them must lie on an even word address (0 mod 2^{19}).

The IBM 370 and the Interdata 8/32 are similar. On both, addresses are measured in bytes; elementary objects must be aligned on a boundary equal to their length, so pointers to short must be 0 mod 2, to int and float 0 mod 4, and to double 0 mod 8. Aggregates are aligned on the strictest boundary required by any of their constituents.

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

$$+ \quad - \quad * \quad / \quad \% \quad \& \quad | \quad \char`^ \quad << \quad >> \quad == \quad != \quad < \quad > \quad <= \quad >=$$

or by the unary operators

$$- \quad \tilde{\ }$$

or by the ternary operator

$$?:$$

Parentheses can be used for grouping, but not for function calls.

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. Portability considerations

Certain parts of C are inherently machine dependent. The following list of potential trouble spots is not meant to be all-inclusive, but to point out the main ones.

Purely hardware issues like word size and the properties of floating point arithmetic and integer division have proven in practice to be not much of a problem. Other facets of the hardware are reflected in differing implementations. Some of these, particularly sign extension (converting a negative character into a negative integer) and the order in which bytes are placed in a word, are a nuisance that must be carefully watched. Most of the others are only minor problems.

The number of register variables that can actually be placed in registers varies from machine to machine, as does the set of valid types. Nonetheless, the compilers all do things properly for their own machine; excess or invalid register declarations are ignored.

Some difficulties arise only when dubious coding practices are used. It is exceedingly unwise to write programs that depend on any of these properties.

The order of evaluation of function arguments is not specified by the language. It is right to left on the PDP-11, and VAX-11, left to right on the others. The order in which side effects take place is also unspecified.

Since character constants are really objects of type int, multi-character character constants may be permitted. The specific implementation is very machine dependent, however, because the order in which characters are assigned to a word varies from one machine to another.

Fields are assigned to words and characters to integers right-to-left on the PDP-11 and VAX-11 and left-to-right on other machines. These differences are invisible to isolated programs which do not indulge in type punning (for example, by converting an int pointer to a char pointer and inspecting the pointed-to storage), but must be accounted for when conforming to externally-imposed storage layouts.

The language accepted by the various compilers differs in minor details. Most notably, the current PDP-11 compiler will not initialize structures containing bit-fields, and does not accept a few assignment operators in certain contexts where the value of the assignment is used.

17. Anachronisms

Since C is an evolving language, certain obsolete constructions may be found in older programs. Although most versions of the compiler support such anachronisms, ultimately they will disappear, leaving only a portability problem behind.

Earlier versions of C used the form *=op* instead of *op=* for assignment operators. This leads to ambiguities, typified by

        x=-1

which actually decrements x since the = and the – are adjacent, but which might easily be intended to assign –1 to x.

The syntax of initializers has changed: previously, the equals sign that introduces an initializer was not present, so instead of

        int x   = 1;

one used

        int x   1;

The change was made because the initialization

        int f   (1+2)

resembles a function declaration closely enough to confuse the compilers.

18. Syntax Summary

This summary of C syntax is intended more for aiding comprehension than as an exact statement of the language.

18.1 Expressions

The basic expressions are:

> *expression:*
> > *primary*
> > * *expression*
> > & *expression*
> > − *expression*
> > ! *expression*
> > ~ *expression*
> > ++ *lvalue*
> > −− *lvalue*
> > *lvalue* ++
> > *lvalue* −−
> > sizeof *expression*
> > (*type-name*) *expression*
> > *expression binop expression*
> > *expression* ? *expression* : *expression*
> > *lvalue asgnop expression*
> > *expression* , *expression*

> *primary:*
> > *identifier*
> > *constant*
> > *string*
> > (*expression*)
> > *primary* (*expression-list$_{opt}$*)
> > *primary* [*expression*]
> > *lvalue* . *identifier*
> > *primary* −> *identifier*

> *lvalue:*
> > *identifier*
> > *primary* [*expression*]
> > *lvalue* . *identifier*
> > *primary* −> *identifier*
> > * *expression*
> > (*lvalue*)

The primary-expression operators

> () [] . −>

have highest priority and group left-to-right. The unary operators

> * & − ! ~ ++ −− sizeof (*type-name*)

have priority below the primary operators but higher than any binary operator, and group right-to-left. Binary operators group left-to-right; they have priority decreasing as indicated below. The conditional operator groups right to left.

binop:

```
*     /     %
+     −
>>    <<
<     >     <=      >=
==    !=
&
^
|
&&
||
?:
```

Assignment operators all have the same priority, and all group right-to-left.

asgnop:

```
=   +=   −=   *=   /=   %=   >>=   <<=   &=   ^=   |=
```

The comma operator has the lowest priority, and groups left-to-right.

18.2 Declarations

declaration:
   decl-specifiers init-declarator-list$_{opt}$;

decl-specifiers:
   type-specifier decl-specifiers$_{opt}$
   sc-specifier decl-specifiers$_{opt}$

sc-specifier:
   auto
   static
   extern
   register
   typedef

type-specifier:
   char
   short
   int
   long
   unsigned
   float
   double
   struct-or-union-specifier
   typedef-name

init-declarator-list:
   init-declarator
   init-declarator , init-declarator-list

init-declarator:
   declarator initializer$_{opt}$

declarator:
   identifier
   (*declarator*)
   * *declarator*
   declarator ( )
   declarator [*constant-expression$_{opt}$*]

struct-or-union-specifier:
        `struct` (*struct-decl-list*)
        `struct` *identifier* (*struct-decl-list*)
        `struct` *identifier*
        `union` (*struct-decl-list*)
        `union` *identifier* (*struct-decl-list*)
        `union` *identifier*

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

struct-declarator:
        declarator
        declarator : *constant-expression*
        : *constant-expression*

initializer:
        = *expression*
        = (*initializer-list*)
        = (*initializer-list ,*)

initializer-list:
        expression
        initializer-list , initializer-list
        (*initializer-list*)

type-name:
        type-specifier abstract-declarator

abstract-declarator:
        empty
        (*abstract-declarator*)
        * *abstract-declarator*
        abstract-declarator ( )
        abstract-declarator [*constant-expression*$_{opt}$]

typedef-name:
        identifier

18.3 Statements

compound-statement:
        (*declaration-list*$_{opt}$ *statement-list*$_{opt}$)

declaration-list:
        declaration
        declaration declaration-list

statement-list:
  statement
  statement statement-list

statement:
  compound-statement
  expression ;
  `if` (*expression*) *statement*
  `if` (*expression*) *statement* `else` *statement*
  `while` (*expression*) *statement*
  `do` *statement* `while` (*expression*) ;
  `for` (*expression-1$_{opt}$* ; *expression-2$_{opt}$* ; *expression-3$_{opt}$*) *statement*
  `switch` (*expression*) *statement*
  `case` *constant-expression* :   *statement*
  `default` : *statement*
  `break` ;
  `continue` ;
  `return` ;
  `return` *expression* ;
  `goto` *identifier* ;
  identifier : *statement*
  ;

18.4 External definitions

program:
  external-definition
  external-definition program

external-definition:
  function-definition
  data-definition

function-definition:
  type-specifier$_{opt}$ function-declarator function-body

function-declarator:
  declarator (*parameter-list$_{opt}$*)

parameter-list:
  identifier
  identifier , *parameter-list*

function-body:
  type-decl-list function-statement

function-statement:
  { *declaration-list$_{opt}$ statement-list* }

data-definition:
  `extern`$_{opt}$ *type-specifier$_{opt}$ init-declarator-list$_{opt}$* ;
  `static`$_{opt}$ *type-specifier$_{opt}$ init-declarator-list$_{opt}$* ;

18.5 Preprocessor

```
#define identifier token-string
#define identifier( identifier , ... , identifier ) token-string
#undef identifier
#include "filename"
#include <filename>
#if constant-expression
#ifdef identifier
#ifndef identifier
#else
#endif
#line constant identifier
```

Recent Changes to C

November 15, 1978

A few extensions have been made to the C language beyond what is described in the reference document ("The C Programming Language," Kernighan and Ritchie, Prentice-Hall, 1978).

1. Structure assignment

Structures may be assigned, passed as arguments to functions, and returned by functions. The types of operands taking part must be the same. Other plausible operators, such as equality comparison, have not been implemented.

There is a subtle defect in the PDP-11 implementation of functions that return structures: if an interrupt occurs during the return sequence, and the same function is called reentrantly during the interrupt, the value returned from the first call may be corrupted. The problem can occur only in the presence of true interrupts, as in an operating system or a user program that makes significant use of signals; ordinary recursive calls are quite safe.

2. Enumeration type

There is a new data type analogous to the scalar types of Pascal. To the type-specifiers in the syntax on p. 193 of the C book add

> *enum-specifier*

with syntax

> *enum-specifier:*
>   enum { *enum-list* }
>   enum *identifier* { *enum-list* }
>   enum *identifier*
>
> *enum-list:*
>   *enumerator*
>   *enum-list , enumerator*
>
> *enumerator:*
>   *identifier*
>   *identifier = constant-expression*

The role of the identifier in the enum-specifier is entirely analogous to that of the structure tag in a struct-specifier; it names a particular enumeration. For example,

```
enum color { chartreuse, burgundy, claret, winedark };
...
enum color *cp, col;
```

makes `color` the enumeration-tag of a type describing various colors, and then declares `cp` as a pointer to an object of that type, and `col` as an object of that type.

The identifiers in the enum-list are declared as constants, and may appear wherever constants are required. If no enumerators with = appear, then the values of the constants begin at 0 and increase by 1 as the declaration is read from left to right. An enumerator with = gives the associated identifier the value indicated; subsequent identifiers continue the progression from the assigned value.

Enumeration tags and constants must all be distinct, and, unlike structure tags and members, are drawn from the same set as ordinary identifiers.

Objects of a given enumeration type are regarded as having a type distinct from objects of all other types, and *lint* flags type mismatches. In the PDP-11 implementation all enumeration variables are treated as if they were `int`.
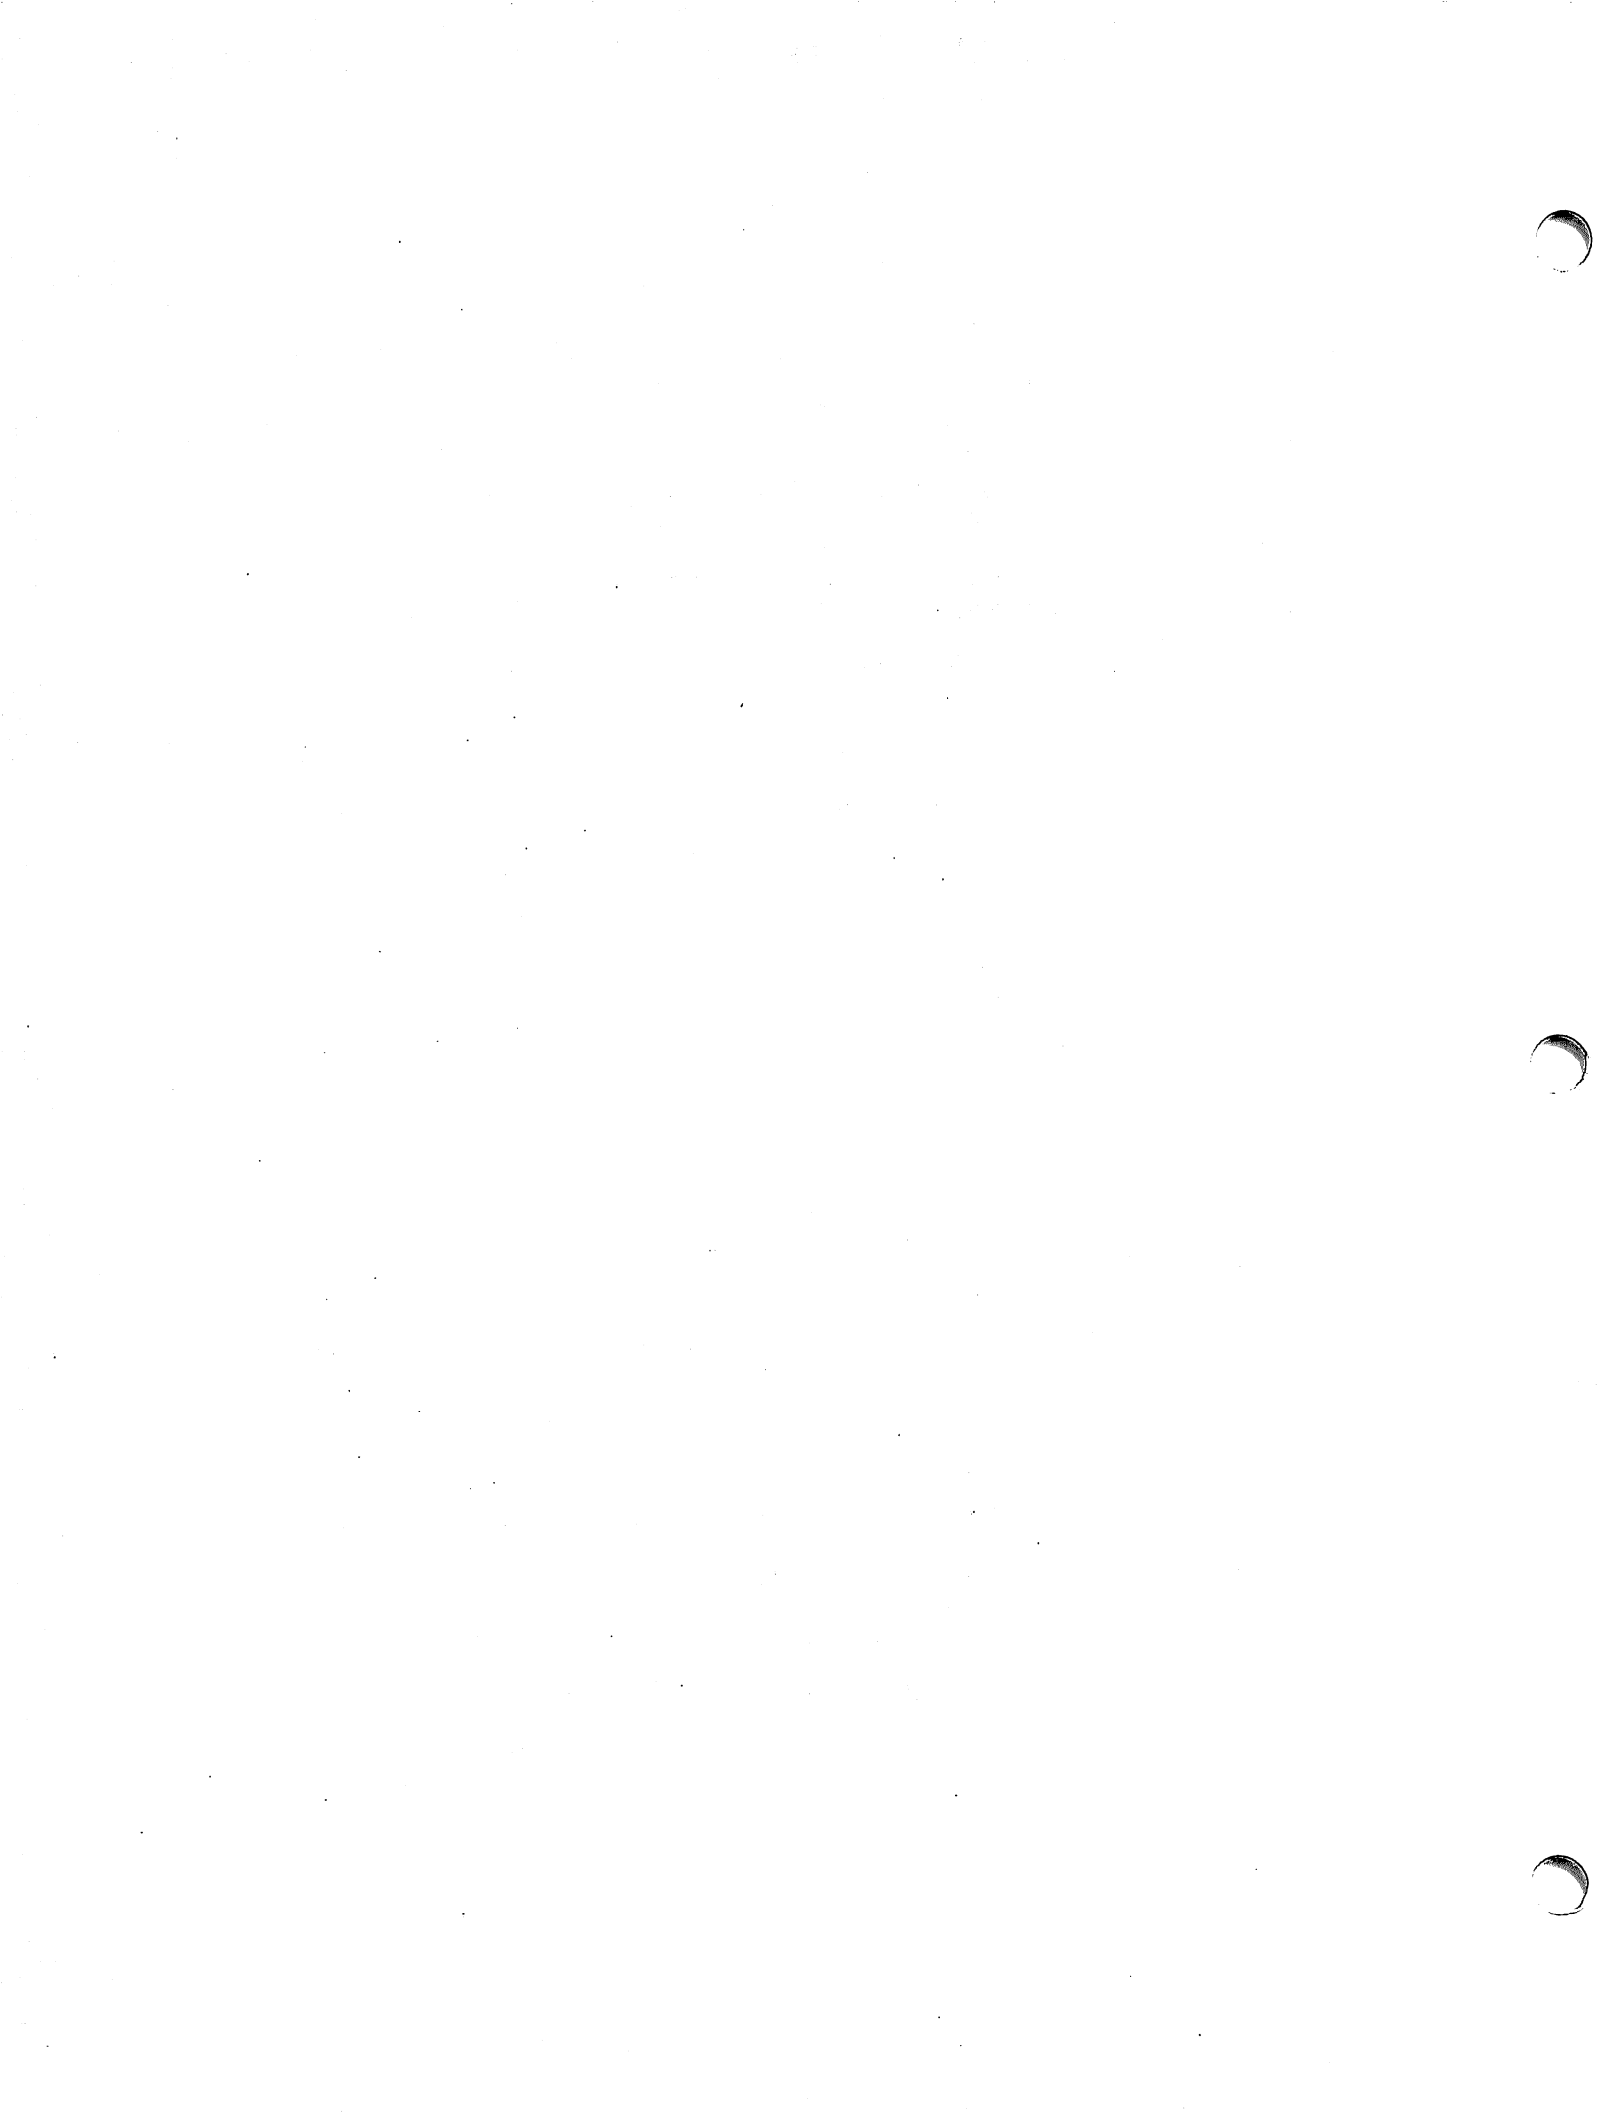
Lint, a C Program Checker

S. C. Johnson

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

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

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

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

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

July 26, 1978

Lint, a C Program Checker

S. C. Johnson

Bell Laboratories
Murray Hill, New Jersey 07974

Introduction and Usage

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

       lint file1.c file2.c

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

       lint —p file1.c file2.c

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

The next several sections describe the major messages; the document closes with sections discussing the implementation and giving suggestions for writing portable C. An appendix gives a summary of the *lint* options.

A Word About Philosophy

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

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

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

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

Unused Variables and Functions

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

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

    extern float sin();

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

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

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

Set/Used Information

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

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

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

Flow of Control

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

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

One form of unreachable statement is not usually complained about by *lint;* a **break** statement that cannot be reached causes no message. Programs generated by *yacc*,[2] and especially *lex*,[3] may have literally hundreds of unreachable **break** statements. The −O flag in the C

compiler will often eliminate the resulting object code inefficiency. Thus, these unreached statements are of little importance, there is typically nothing the user can do about them, and the resulting messages would clutter up the *lint* output. If these messages are desired, *lint* can be invoked with the **−b** option.

Function Values

Sometimes functions return values which are never used; sometimes programs incorrectly use function "values" which have never been returned. *Lint* addresses this problem in a number of ways.

Locally, within a function definition, the appearance of both

        return(*expr*);

and

        return ;

statements is cause for alarm; *lint* will give the message

        function *name* contains return(e) and return

The most serious difficulty with this is detecting when a function return is implied by flow of control reaching the end of the function. This can be seen with a simple example:

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

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

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

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

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

Type Checking

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

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

The type checking rules also require that, in structure references, the left operand of the −> be a pointer to structure, the left operand of the . be a structure, and the right operand of

these operators be a member of the structure implied by the left operand. Similar checking is done for references to unions.

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

With enumerations, checks are made that enumeration variables or members are not mixed with other types, or other enumerations, and that the only operations applied are =, initialization, ==, !=, and function arguments and return values.

Type Casts

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

        p = 1 ;

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

        p = (char *)1 ;

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

Nonportable Character Use

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

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

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

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

Assignments of longs to ints

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

Strange Constructions

Several perfectly legal, but somewhat strange, constructions are flagged by *lint;* the messages hopefully encourage better code quality, clearer style, and may even point out bugs. The —h flag is used to enable these checks. For example, in the statement

        *p++ ;

the * does nothing; this provokes the message "null effect" from *lint.* The program fragment

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

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

        if(x > 0) ...

is equivalent to

        if(x != 0)

which may not be the intended action. *Lint* will say "degenerate unsigned comparison" in these cases. If one says

        if(1 != 0) ....

lint will report "constant in conditional context", since the comparison of 1 with 0 gives a constant result.

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

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

or

        x<<2 + 40

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

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

Ancient History

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

The older forms of assignment operators (e.g., =+, =−, . . .) could cause ambiguous expressions, such as

        a =−1 ;

which could be taken as either

        a =− 1 ;

or

        a = −1 ;

The situation is especially perplexing if this kind of ambiguity arises as the result of a macro substitution. The newer, and preferred operators (+=, −=, etc.) have no such ambiguities. To spur the abandonment of the older forms, *lint* complains about these old fashioned

operators.

A similar issue arises with initialization. The older language allowed

    int x 1 ;

to initialize x to 1. This also caused syntactic difficulties: for example,

    int x (−1) ;

looks somewhat like the beginning of a function declaration:

    int x (y) { . . .

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

    int x = −1 ;

This is free of any possible syntactic ambiguity.

Pointer Alignment

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

Multiple Uses and Side Effects

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

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

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

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

will draw the complaint:

    warning: i evaluation order undefined

Implementation

Lint consists of two programs and a driver. The first program is a version of the Portable
C Compiler[4, 5] which is the basis of the IBM 370, Honeywell 6000, and Interdata 8/32 C com-
pilers. This compiler does lexical and syntax analysis on the input text, constructs and main-
tains symbol tables, and builds trees for expressions. Instead of writing an intermediate file

which is passed to a code generator, as the other compilers do, *lint* produces an intermediate file which consists of lines of ascii text. Each line contains an external variable name, an encoding of the context in which it was seen (use, definition, declaration, etc.), a type specifier, and a source file name and line number. The information about variables local to a function or file is collected by accessing the symbol table, and examining the expression trees.

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

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

Portability

C on the Honeywell and IBM systems is used, in part, to write system code for the host operating system. This means that the implementation of C tends to follow local conventions rather than adhere strictly to UNIX† system conventions. Despite these differences, many C programs have been successfully moved to GCOS and the various IBM installations with little effort. This section describes some of the differences between the implementations, and discusses the *lint* features which encourage portability.

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

        int a ;

outside of any function. The UNIX loader will resolve these declarations, and cause only a single word of storage to be set aside for *a*. Under the GCOS and IBM implementations, this is not feasible (for various stupid reasons!) so each such declaration causes a word of storage to be set aside and called *a*. When loading or library editing takes place, this causes fatal conflicts which prevent the proper operation of the program. If *lint* is invoked with the −p flag, it will detect such multiple definitions.

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

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

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

---

†UNIX is a Trademark of Bell Laboratories.

x &= 0177700 ;

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

x &= ~ 077 ;

which will work on all these machines.

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

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

Shutting Lint Up

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

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

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

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

/* NOTREACHED */

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

/* NOSTRICT */

can be used; the situation reverts to the previous default after the next expression. The −v flag can be turned on for one function by the directive

/* ARGSUSED */

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

```
/* VARARGS */
```

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

```
/* VARARGS2 */
```

will cause the first two arguments to be checked, the others unchecked. Finally, the directive

```
/* LINTLIBRARY */
```

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

Library Declaration Files

Lint accepts certain library directives, such as

```
-ly
```

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

```
/* LINTLIBRARY */
```

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

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

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

Bugs, etc.

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

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

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

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

pressures to add even more of these options.

In conclusion, it appears that the general notion of having two programs is a good one. The compiler concentrates on quickly and accurately turning the program text into bits which can be run; *lint* concentrates on issues of portability, style, and efficiency. *Lint* can afford to be wrong, since incorrectness and over-conservatism are merely annoying, not fatal. The compiler can be fast since it knows that *lint* will cover its flanks. Finally, the programmer can concentrate at one stage of the programming process solely on the algorithms, data structures, and correctness of the program, and then later retrofit, with the aid of *lint*, the desirable properties of universality and portability.

References

1. B. W. Kernighan and D. M. Ritchie, *The C Programming Language,* Prentice-Hall, Englewood Cliffs, New Jersey (1978).

2. S. C. Johnson, "Yacc — Yet Another Compiler-Compiler," Comp. Sci. Tech. Rep. No. 32, Bell Laboratories, Murray Hill, New Jersey (July 1975).

3. M. E. Lesk, "Lex — A Lexical Analyzer Generator," Comp. Sci. Tech. Rep. No. 39, Bell Laboratories, Murray Hill, New Jersey (October 1975).

4. S. C. Johnson and D. M. Ritchie, "UNIX Time-Sharing System: Portability of C Programs and the UNIX System," *Bell Sys. Tech. J.* 57(6) pp. 2021-2048 (1978).

5. S. C. Johnson, "A Portable Compiler: Theory and Practice," *Proc. 5th ACM Symp. on Principles of Programming Languages,* (January 1978).

Appendix: Current Lint Options

The command currently has the form

lint [−options] files... library-descriptors...

The options are

h   Perform heuristic checks

p   Perform portability checks

v   Don't report unused arguments

u   Don't report unused or undefined externals

b   Report unreachable **break** statements.

x   Report unused external declarations

a   Report assignments of **long** to **int** or shorter.

c   Complain about questionable casts

n   No library checking is done

s   Same as **h** (for historical reasons)
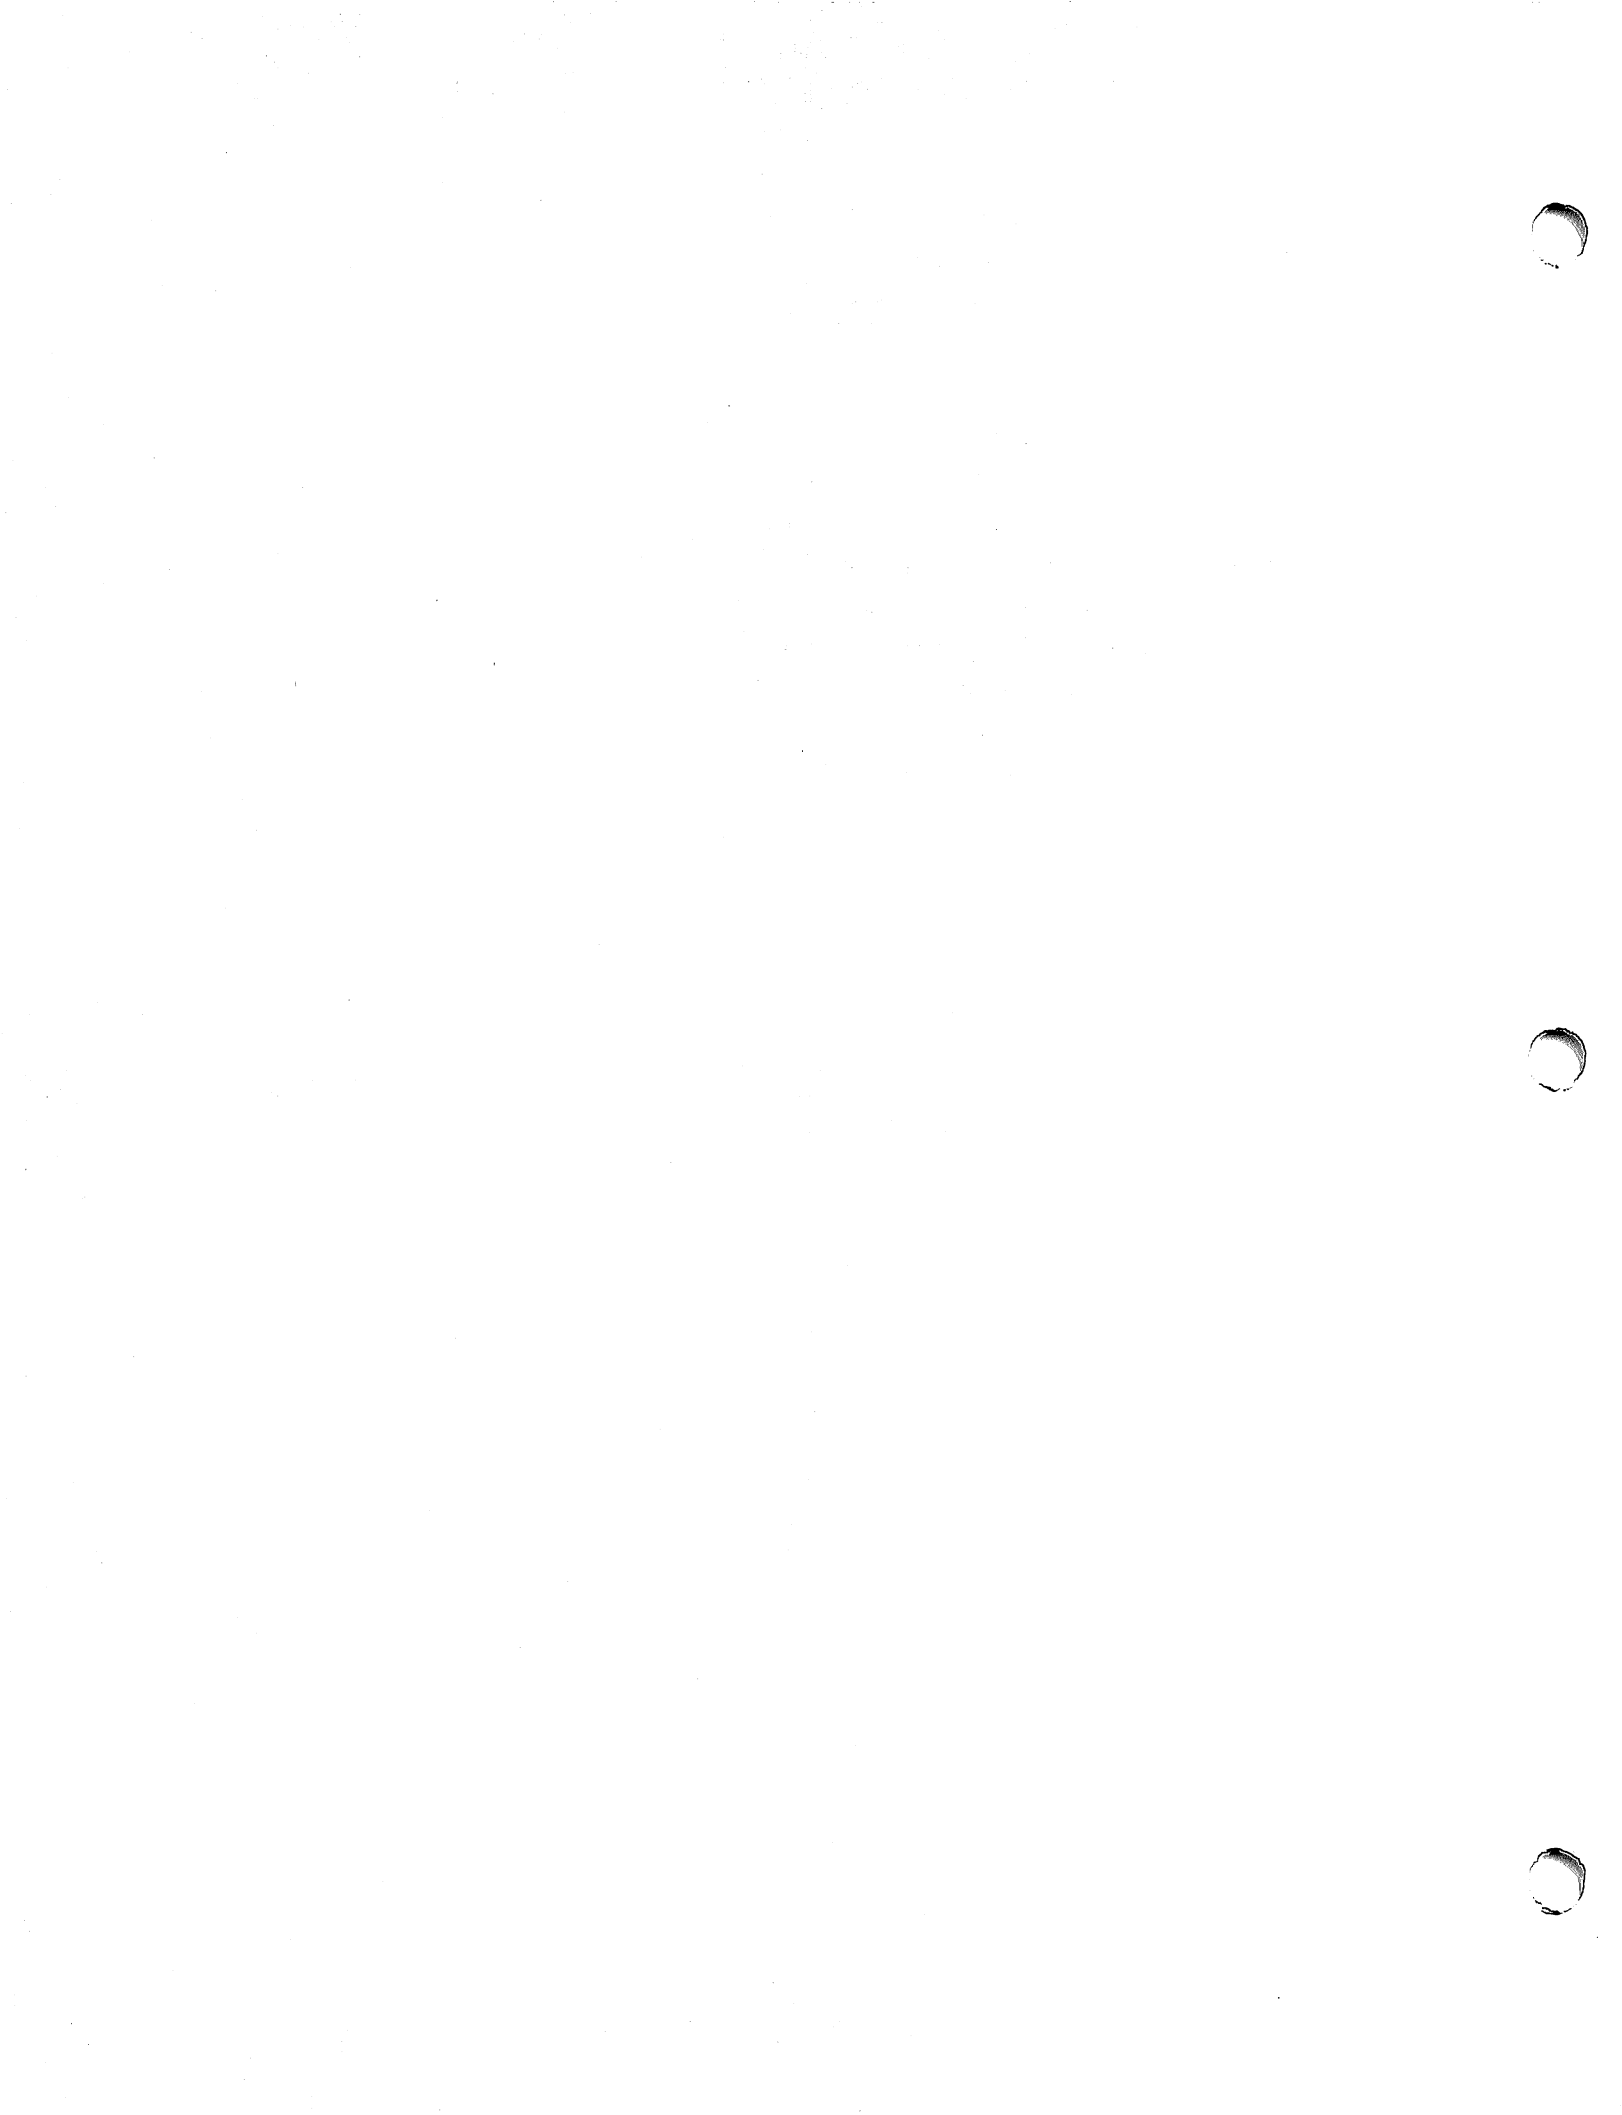
ADDENDUM

Beginning with the UNIX^{TM} System III release, the user interface with <u>lint</u> has changed. In the past, some command line parameters were used to suppress <u>lint</u> checks; other parameters enabled associated checks. In the current <u>lint</u>, all parameters suppress associated checks. By default, all the checks are made.

The output format of <u>lint</u> has also been changed. Most error messages are grouped together by type, and the format has been redesigned to be more readable.

Make — A Program for Maintaining Computer Programs

S. I. Feldman

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

In a programming project, it is easy to lose track of which files need to be reprocessed or recompiled after a change is made in some part of the source. *Make* provides a simple mechanism for maintaining up-to-date versions of programs that result from many operations on a number of files. It is possible to tell *Make* the sequence of commands that create certain files, and the list of files that require other files to be current before the operations can be done. Whenever a change is made in any part of the program, the *Make* command will create the proper files simply, correctly, and with a minimum amount of effort.

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

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

August 15, 1978

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  -lS  -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  -lS  -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 −lS"

loads them with both the Lex (``−ll'') and the Standard (``−lS'') 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 is 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:

---

†UNIX is a Trademark of Bell Laboratories.

```
2 = xyz
abc = −ll −ly −lS
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" flags 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., *cd* 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:

.o	Object file
.c	C source file
.e	Efl source file
.r	Ratfor source file
.f	Fortran source file
.s	Assembler source file
.y	Yacc-C source grammar
.yr	Yacc-Ratfor source grammar
.ye	Yacc-Efl source grammar
.l	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.
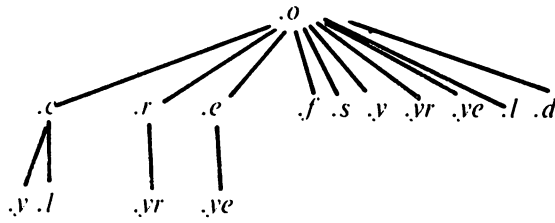
If the file *x.o* were needed and there were an *x.c* in the description or directory, it would be compiled. If there were also an *x.l*, that grammar would be run through Lex before compiling the result. However, if there were no *x.c* but there were an *x.l*, *make* would discard the intermediate C-language 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= −lS
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. C. Johnson, "Yacc — Yet Another Compiler-Compiler", Bell Laboratories Computing Science Technical Report #32, July 1978.

2.   M. E. Lesk, "Lex — A Lexical Analyzer Generator", Computing Science Technical Report #39, October 1975.

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 $@
```

An Augmented Version of MAKE

E. G. Bradford

Bell Laboratories
Whippany, New Jersey 07981

ABSTRACT

This paper describes an augmented version of the UNIX† *make* command. With one debatable exception, this version is completely upward compatible with the old version. This paper describes and gives examples only of additional features. The reader is assumed to have read the original *make* paper by S. I. Feldman.[1] Further possible developments for *make* are also discussed.

1. INTRODUCTION

This paper describes in some detail an augmented version of the *make* program. Some justification will be given for the chosen implementation and examples will demonstrate the additional features.

2. MOTIVATION FOR THE CURRENT IMPLEMENTATION

The *make* program was originally written for personal use by S. I. Feldman. However, it became popular on the BTL Research UNIX machine and a more formal version was built and installed for general use. Further developments of *make* have not been necessary in the BTL Research environment, and thus have not been done.

Elsewhere, *make* was perceived as an excellent program administrative tool and has been used extensively in at least one project for over two years. However, *make* had many shortcomings: handling of libraries was tedious; handling of the SCCS[2] file-name format was difficult or impossible; environment variables were completely ignored by *make*, and the general lack of ability to maintain files in a remote directory. These shortcomings hindered large scale use of *make* as a program support tool.

Make has been modified to handle the problems above. The additional features are within the original syntactic framework of *make* and few if any new syntactical entities have been introduced. A notable exception is the *include* file capability. Further, most of the additions result in a "Don't know how to make ..." message from the old version of *make*.

3. THE ADDITIONAL FEATURES

The following paragraphs describe with examples the additional features of the *make* program. In general, the examples are taken from existing *makefiles*. Also, the appendices are working *makefiles*.

3.1 The Environment Variables

Environment variables are read and added to the macro definitions each time *make* executes. Precedence is a prime consideration in doing this properly. For example, if the environment variable CC is set to occ, does it override the command line? Does it override the definition in the makefile? The following describes *make*'s interaction with the environment.

---

† UNIX is a trademark of Bell Laboratories.

1. *MAKE – A Program for Maintaining Computer Programs* by S. I. Feldman.

2. *Source Code Control System User's Guide* by L. E. Bonanni and C. A. Salemi.

A new macro, MAKEFLAGS is maintained by *make*. It is defined as the collection of all input flag arguments into a string (without minus signs). It is exported, and thus accessible to further invocations of *make*. Command line flags and assignments in the *makefile* update MAKEFLAGS. Thus, to describe how the environment interacts with *make*, we also need to consider the MAKEFLAGS macro (environment variable).

When executed, *make* assigns macro definitions in the following order:

1. Read the MAKEFLAGS environment variable. If it is not present or null, the internal *make* variable MAKEFLAGS is set to the null string. Otherwise, each letter in MAKEFLAGS is assumed to be an input flag argument and is processed as such. (The only exceptions are the −f, −p, and −r flags.)
2. Read and set the input flags from the command line. The command line adds to the previous settings from the MAKEFLAGS environment variable.
3. Read macro definitions from the command line. These are made *not resettable*. Thus any further assignments to these names are ignored.
4. Read the internal list of macro definitions. These are found in the file *rules.c* of the source for *make*. (See Appendix A for the complete makefile which represents the internally defined macros and rules.) They give default definitions for the C compiler (CC=cc), the assembler (AS=as), etc.
5. Read the environment. The environment variables are treated as macro definitions and marked as *exported* (in the shell sense). Note, MAKEFLAGS will be read and set again. However, because it is not an internally defined variable (in *rules.c*), this has the effect of doing the same assignment twice. The exception to this is when MAKEFLAGS is assigned on the command line. (The reason it was read previously, was to be able to turn the debug flag on before anything else was done.)
6. Read the *makefile*(s). The assignments in the *makefile*(s) will override the environment. This order was chosen so when one reads a makefile and executes *make* one knows what to expect. That is, one gets what one sees unless the −e flag is used. The −e is an additional command-line flag that tells *make* to have the environment override the *makefile* assignments. Thus if make −e ... is typed, the variables in the environment override the definitions in the *makefile*. (Note, there is no way to override the command line assignments.) Also note that if MAKEFLAGS is assigned it will override the environment. (This would be useful for further invocations of *make* from the current *makefile*.)

It may be clearer to list the precedence of assignments. Thus, in order from least binding to most binding, we have:

1. Internal definitions (from *rules.c*).
2. Environment.
3. *Makefile*(s).
4. Command line.

The −e flag has the effect of changing the order to:

1. Internal definitions (from *rules.c*).
2. *Makefile*(s).
3. Environment.
4. Command line.

This ordering is general enough to allow a programmer to define a *makefile* or set of *makefiles* whose parameters are dynamically definable.

3.2 Recursive Makefiles

Another feature was added to *make* concerning the environment and recursive invocations. If the sequence $(MAKE) appears anywhere in a shell command line, the line will be executed even if the −n flag is set. Because the −n flag is exported across invocations of *make*,

(through the MAKEFLAGS variable) the only thing which will actually get executed is the *make* command itself. This feature is useful when a hierarchy of *makefile(s)* describes a set of software subsystems. For testing purposes, make —n ... can be executed and everything that would have been done will get printed out; including output from lower level invocations of *make*.

3.3 Format of Shell Commands within Make

Make remembers embedded new-lines and tabs in shell command sequences. Thus, if the programmer puts a *for* loop in the makefile with indentation, when *make* prints it out, it retains the indentation and backslashes. The output can still be piped to the shell and is readable. This is obviously a cosmetic change; no new functionality is gained.

3.4 Archive Libraries

Make has an improved interface to archive libraries. Due to a lack of documentation, most people are probably not aware of the current syntax for addressing members of archive libraries. The previous version of *make* allows a user to name a member of a library in the following manner:

lib(object.o)

or:

lib((_localtime))

where the second method actually refers to an entry point of an object file within the library. (*Make* looks through the library, locates the entry point and translates it to the correct object file name.)

To use this procedure to maintain an archive library, the following type of *makefile* is required:

```
lib::    lib(ctime.o)
         $(CC) —c —O ctime.c
         ar rv lib ctime.o
         rm ctime.o
lib::    lib(fopen.o)
         $(CC) —c —O fopen.c
         ar rv lib fopen.o
         rm fopen.o
... and so on for each object ...
```

This is tedious and error prone. Obviously, the command sequences for adding a C file to a library are the same for each invocation, the file name being the only difference each time. (This is true in most cases.) Similarly for assembler and YACC and LEX files.

The current version gives the user access to a rule for building libraries. The *handle* for the rule is the .a suffix. Thus a .c.a rule is the rule for compiling a C source file, adding it to the library, and removing the .o cadaver. Similarly, the .y.a, the .s.a and the .l.a rules rebuild YACC, assembler, and LEX files, respectively. The current archive rules defined internally are .c.a, .c~.a, and .s~.a. (The tilde (~) syntax will be described shortly.) The user may define in his makefile any other rules he may need.

The above two-member library is then maintained with the following shorter makefile:

```
lib:     lib(ctime.o) lib(fopen.o)
         @echo lib up-to-date.
```

The internal rules are already defined to complete the preceding library maintenance. The actual .c.a rules is as follows:

```
.c.a:
        $(CC)  -c  $(CFLAGS)  $<
        ar rv $@ $*.o
        rm  -f $*.o
```

Thus, the $@ macro is the .a target (lib) and the $< and $* macros are set to the out-of-date
C file and the file name without suffix, respectively (*ctime.c* and *ctime*). The $< macro (in the
preceding rule) could have been changed to $*.c.

It might be useful to go into some detail about exactly what *make* does when it sees the con-
struction:

```
lib:     lib(ctime.o)
         @echo lib up-to-date
```

Assume the object in the library is out-of-date with respect to *ctime.c*. Also, there is no *ctime.o*
file:

1. Do *lib*.
2. To do *lib*, do each dependent of *lib*.
3. Do *lib(ctime.o)*.
4. To do *lib(ctime.o)*, do each dependent of *lib(ctime.o)*. (There are none.)
5. Use internal rules to try to build *lib(ctime.o)*. (There is no explicit rule.) Note that
    lib(ctime.o) has a parenthesis in the name so identify the target suffix as .a. (This is the
    key. There is no explicit .a at the end of the *lib* library name. The parenthesis forces the
    .a suffix.) In this sense, the .a is hard-wired into *make*.
6. Break the name *lib(ctime.o)* up into *lib* and *ctime.o*. Define two macros, $@ (=*lib*) and $*
    (=*ctime*).
7. Look for a rule .X.a and a file $*.X. The first .X (in the .SUFFIXES list) which fulfills
    these conditions is .c so the rule is .c.a and the file is *ctime.c*. Set $< to be *ctime.c* and
    execute the rule. (In fact, *make* must then do *ctime.c*. However, the search of the current
    directory yields no other candidates, whence, the search ends.)
8. The library has been updated. Do the rule associated with the lib: dependency; namely:

        echo lib up-to-date

It should be noted that to let *ctime.o* have dependencies, the following syntax is required:

    lib(ctime.o): $(INCDIR)/stdio.h

Thus, explicit references to .o files are unnecessary. There is also a new macro for referencing
the archive member name when this form is used. $% is evaluated each time $@ is evaluated.
If there is no current archive member, $% is null. If an archive member exists, then $% evalu-
ates to the expression between the parentheses.

An example *makefile* for a larger library is given in Appendix B. The reader will note also that
there are no lingering *.o files left around. The result is a library maintained directly from the
source files (or more generally, from the SCCS files).

3.5 SCCS File Names: The Tilde

The syntax of *make* does not directly permit referencing of prefixes. For most types of files on
UNIX machines this is acceptable, because nearly everyone uses a suffix to distinguish different
types of files. SCCS files are the exception. Here, s. precedes the file-name part of the com-
plete path name.

To allow *make* easy access to the prefix s. requires either a redefinition of the rule naming syn-
tax of *make* or a trick. The trick is to use the tilde (~) as an identifier of SCCS files. Hence,
.c~.o refers to the rule which transforms an SCCS C source file into an object. Specifically, the
internal rule is.

```
.c~.o:
        $(GET) $(GFLAGS) -p $< > $*.c
        $(CC) $(CFLAGS) -c $*.c
        -rm -f $*.c
```

Thus the tilde appended to any suffix transforms the file search into an SCCS file-name search with the actual suffix named by the dot and all characters up to (but not including) the tilde. The following SCCS suffixes are internally defined:

```
.c~
.y~
.s~
.sh~
.h~
```

The following rules involving SCCS transformations are internally defined:

```
.c~:
.sh~:
.c~.o:
.s~.o:
.y~.o:
.l~.o:
.y~.c:
.c~.a:
.s~.a:
.h~.h:
```

Obviously, the user can define other rules and suffixes which may prove useful. The tilde gives him a handle on the SCCS file-name format so that this is possible.

3.6 The Null Suffix

In the UNIX source code, there are many commands that consist of a single source file. It seemed wasteful to maintain an object of such files for *make*'s pleasure. The current implementation supports single-suffix rules, or, if one prefers, a null suffix. Thus, to maintain the program *cat* one needs a rule in the *makefile* of the following form:

```
.c:
        $(CC) -n -O $< -o $@
```

In fact, this .c: rule is internally defined so no *makefile* is necessary at all! One need only type:

```
make cat dd echo date
```

(these are notable single-file programs) and all four C source files are passed through the above shell command line associated with the .c: rule. The internally defined single-suffix rules are:

```
.c:
.c~:
.sh:
.sh~:
```

Others may be added in the *makefile* by the user.

3.7 Include Files

Make has an include file capability. If the string Include appears as the first seven letters of a line in a *makefile* and is followed by a blank or a tab, the following string is assumed to be a file name that the current invocation of *make* will read. The file descriptors are stacked for reading *include* files so no more that about sixteen levels of nested includes is supported.

3.8 Invisible SCCS Makefiles

SCCS *makefiles* are invisible to *make*. That is, if *make* is typed and only a file named

s.makefile exists, *make* will do a *get*(1) on it, then read it and remove it; likewise for —f arguments and *include* files.

3.9 Dynamic Dependency Parameters

A new dependency parameter has been defined. It has meaning only on the dependency line in a makefile. $$@ refers to the current "thing" to the left of the colon (which is $@). Also the form $$(@F) exists which allows access to the file part of $@. Thus, in the following:

    cat: $$@.c

the dependency is translated at execution time to the string cat.c. This is useful for building large numbers of executable files, each of which has only one source file. For instance the UNIX command directory could have a *makefile* like:

    CMDS = cat dd echo date cc cmp comm ar ld chown

    $(CMDS): $$@.c
            $(CC) —O $? —o $@

Obviously, this is a subset of all the single-file programs. For multiple-file programs, a directory is usually allocated and a separate *makefile* is made. For any particular file which has a peculiar compilation procedure, a specific entry must be made in the *makefile*.

The second useful form of the dependency parameter is $$(@F). It represents the file-name part of $$@. Again, it is evaluated at execution time. Its usefulness becomes evident when trying to maintain the /usr/include directory from a makefile in the /usr/src/head directory. Thus the /usr/src/head/makefile would look like:

    INCDIR = /usr/include

    INCLUDES = \
            $(INCDIR)/stdio.h \
            $(INCDIR)/pwd.h \
            $(INCIDR)/dir.h \
            $(INCDIR)/a.out.h

    $(INCLUDES): $$(@F)
            cp $? $@
            chmod 0444 $@

This would completely maintain the /usr/include directory whenever one of the above files in /usr/src/head was updated.

3.10 Extensions of $*, $@, and $<

The internally generated macros $*, $@, and $< are useful generic terms for current targets and out-of-date relatives. To this list has been added the following related macros: $(@D), $(@F), $(*D), $(*F), $(<D), and $(<F). The D refers to the directory part of the single-letter macro. The F refers to the file-name part of the single-letter macro. These additions are useful when building hierarchical makefiles. They allow access to directory names for purposes of using the *cd* command of the shell. Thus, a shell command can be:

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

An interesting example of the use of these features can be found in the set of *makefiles* in Appendix C.

3.11 Output Translations

Macros in shell commands can now be translated when evaluated. The form is as follows:

$(macro:string1 = string2)

The meaning is that $(macro) is evaluated. For each appearance of *string1* in the evaluated macro, *string2* is substituted. The meaning of finding *string1* in $(macro) is that the evaluated $(macro) is considered as a bunch of strings each delimited by white space (blanks or tabs). Thus the occurrence of *string1* in $(macro) means that a regular expression of the following form has been found:

.*<string1>[TAB|BLANK]

This particular form was chosen because *make* usually concerns itself with suffixes. A more general regular expression match could be implemented if the need arises. The usefulness of this type of translation occurs when maintaining archive libraries. Now, all that is necessary is to accumulate the out-of-date members and write a shell script which can handle all the C programs (i.e., those file ending in .c). Thus the following fragment will optimize the executions of make for maintaining an archive library:

```
$(LIB): $(LIB)(a.o)  $(LIB)(b.o)  $(LIB)c.o)
        $(CC)  -c  $(CFLAGS)  $(?:.o=.c)
        ar  rv  $(LIB)  $?
        rm  $?
```

A dependency of the preceding form would be necessary for each of the different types of source files (suffices) which define the archive library. These translations are added in an effort to make more general use of the wealth of information that *make* generates.

4. CONCLUSIONS

The augmentations described above have increased the size of *make* significantly, but it is our belief that this increase in size is a reasonable price to pay for the resulting additional features.

Appendix A: Internal Definitions

The following *makefile* will exactly reproduce the internal rules of the current version of *make*. Thus if **make** —r ... is typed and a *makefile* includes this *makefile* the results would be identical to excluding the —r option and the *include* line in the makefile. Note that this output can be reproduced by:

    make —fp — < /dev/null 2>/dev/null

The output will appear on the standard output.)

```
#       LIST OF SUFFIXES
.SUFFIXES: .o .c .c~ .y .y~ .l .l~ .s .s~ .sh .sh~ .h .h~

#       PRESET VARIABLES
MAKE=make
YACC=yacc
YFLAGS=
LEX=lex
LFLAGS=
LD=ld
LDFLAGS=
CC=cc
CFLAGS=—O
AS=as
ASFLAGS=
GET=get
GFLAGS=

#       SINGLE SUFFIX RULES
.c:
        $(CC) —n —O $< —o $@
.c~:
        $(GET) $(GFLAGS) —p $< > $*.c
        $(CC) —n —O $*.c —o $*
        —rm —f $*.c
.sh:
        cp $< $@
.sh~:
        $(GET) $(GFLAGS) —p $< > $*.sh
        cp $*.sh $*
        —rm —f $*.sh

#       DOUBLE SUFFIX RULES
.c.o:
        $(CC) $(CFLAGS) —c $<
.c~.o:
        $(GET) $(GFLAGS) —p $< > $*.c
        $(CC) $(CFLAGS) —c $*.c
        —rm —f $*.c
.c~.c:
        $(GET) $(GFLAGS) —p $< > $*.c
.s.o:
        $(AS) $(ASFLAGS) —o $@ $<
```

```
.s~.o:
        $(GET) $(GFLAGS) -p $< > $*.s
        $(AS) $(ASFLAGS) -o $*.o $*.s
        -rm -f $*.s
.y.o:
        $(YACC) $(YFLAGS) $<
        $(CC) $(CFLAGS) -c y.tab.c
        rm y.tab.c
        mv y.tab.o $@
.y~.o:
        $(GET) $(GFLAGS) -p $< > $*.y
        $(YACC) $(YFLAGS) $*.y
        $(CC) $(CFLAGS) -c y.tab.c
        rm -f y.tab.c $*.y
        mv y.tab.o $*.o
.l.o:
        $(LEX) $(LFLAGS) $<
        $(CC) $(CFLAGS) -c lex.yy.c
        rm lex.yy.c
        mv lex.yy.o $@
.l~.o:
        $(GET) $(GFLAGS) -p $< > $*.l
        $(LEX) $(LFLAGS) $*.l
        $(CC) $(CFLAGS) -c lex.yy.c
        rm -f lex.yy.c $*.l
        mv lex.yy.o $*.o
.y.c :
        $(YACC) $(YFLAGS) $<
        mv y.tab.c $@
.y~.c:
        $(GET) $(GFLAGS) -p $< > $*.y
        $(YACC) $(YFLAGS) $*.y
        mv y.tab.c $*.c
        -rm -f $*.y
.l.c :
        $(LEX) $<
        mv lex.yy.c $@
.c.a:
        $(CC) -c $(CFLAGS) $<
        ar rv $@ $*.o
        rm -f $*.o
.c~.a:
        $(GET) $(GFLAGS) -p $< > $*.c
        $(CC) -c $(CFLAGS) $*.c
        ar rv $@ $*.o
        rm -f $*.[co]
.s~.a:
        $(GET) $(GFLAGS) -p $< > $*.s
        $(AS) $(ASFLAGS) -o $*.o $*.s
        ar rv $@ $*.o
        -rm -f $*.[so]
.h~.h:
        $(GET) $(GFLAGS) -p $< > $*.h
```

Appendix B: A Library Makefile

The following library maintaining makefile is from current work on LSX. It completely maintains the LSX operating system library.

```
#       @(#)/usr/src/cmd/make/make.tm       3.2
LIB  =  lsxlib

PR  =  vpr  −b LSX

INSDIR  =  /r1/flop0/
INS  =  eval

lsx::   $(LIB) low.o mch.o
        ld  −x low.o mch.o $(LIB)
        mv a.out lsx
        @size lsx  ·

#       Here, $(INS) is either : or eval.
lsx::
        $(INS) 'cp lsx $(INSDIR)lsx && \
                strip $(INSDIR)lsx && \
                ls  −l $(INSDIR)lsx'

print:
        $(PR) header.s low.s mch.s •.h •.c Makefile

$(LIB): \
        $(LIB)(clock.o) \
        $(LIB)(main.o) \
        $(LIB)(tty.o) \
        $(LIB)(trap.o) \
        $(LIB)(sysent.o) \
        $(LIB)(sys2.o) \
        $(LIB)(sys3.o) \
        $(LIB)(sys4.o) \
        $(LIB)(sys1.o) \
        $(LIB)(sig.o) \
        $(LIB)(fio.o) \
        $(LIB)(kl.o) \
        $(LIB)(alloc.o) \
        $(LIB)(nami.o) \
        $(LIB)(iget.o) \
        $(LIB)(rdwri.o) \
        $(LIB)(subr.o) \
        $(LIB)(bio.o) \
        $(LIB)(decfd.o) \
        $(LIB)(slp.o) \
        $(LIB)(space.o) \
        $(LIB)(puts.o)
        @echo $(LIB) now up-to-date.

.s.o:
        as  −o $•.o header.s $•.s
.o.a:
        ar rv $@ $<
        rm  −f $<
```

```
.s.a:
        as —o $•.o header.s $•.s
        ar rv $@ $•.o
        rm —f $•.o

.PRECIOUS:    $(LIB)
```

Appendix C: Example of Recursive Use of Makefiles

The following set of *makefiles* maintain the operating system for Columbus UNIX. They are included here to provide realistic examples of the use of *make*. Each one is named 70.mk. The following command forces a complete rebuild of the operating system:

    FRC=FRC make —f 70.mk

where the current directory is ucb. Here, we have used some of the conventions described in A. Chellis's paper entitled *Proposed Structure for UNIX/TS and UNIX/RT Makefiles*. FRC is a convention for FoRCing *make* to completely rebuild a target starting from scratch.

./ucb makefile

```
#       @(#)/usr/src/cmd/make/make.tm    3.2
#       ucb/70.mk makefile

VERSION = 70

DEPS = \
        os/low.$(VERSION).o \
        os/mch.$(VERSION).o \
        os/conf.$(VERSION).o \
        os/lib1.$(VERSION).a \
        io/lib2.$(VERSION).a

#       This makefile will re-load unix.$(VERSION) if any
#       of the $(DEPS) is out-of-date wrt unix.$(VERSION).
#       Note, it will not go out and check each member
#       of the libraries.  To do this, the FRC macro must
#       be defined.

unix.$(VERSION):    $(DEPS) $(FRC)
        load —s $(VERSION)

$(DEPS):        $(FRC)
        cd $(@D); $(MAKE) —f $(VERSION).mk $(@F)

all:    unix.$(VERSION)
        @echo unix.$(VERSION) up-to-date.

includes:
        cd head/sys; $(MAKE) —f $(VERSION).mk

FRC:    includes;
```

```
#        @ (#)/usr/src/cmd/make/make.tm        3.2
#        ucb/os/70.mk makefile
VERSION = 70

LIB = lib1.$(VERSION).a
COMPOOL=

LIBOBJS = \
        $(LIB)(main.o) \
        $(LIB)(alloc.o) \
        $(LIB)(iget.o) \
        $(LIB)(prf.o) \
        $(LIB)(rdwri.o) \
        $(LIB)(slp.o) \
        $(LIB)(subr.o) \
        $(LIB)(text.o) \
        $(LIB)(trap.o) \
        $(LIB)(sig.o) \
        $(LIB)(sysent.o) \
        $(LIB)(sys1.o) \
        $(LIB)(sys2.o) \
        $(LIB)(sys3.o) \
        $(LIB)(sys4.o) \
        $(LIB)(sys5.o) \
        $(LIB)(syscb.o) \
        $(LIB)(maus.o) \
        $(LIB)(messag.o) \
        $(LIB)(nami.o) \
        $(LIB)(fio.o) \
        $(LIB)(clock.o) \
        $(LIB)(acct.o) \
        $(LIB)(errlog.o)

ALL = \
        conf.$(VERSION).o \
        low.$(VERSION).o \
        mch.$(VERSION).o \
        $(LIB)

all:    $(ALL)
        @echo `$(ALL)' now up-to-date.

$(LIB)::$(LIBOBJS)

$(LIBOBJS):    $(FRC);

FRC:
        rm -f $(LIB)

clobber:cleanup
        -rm -f $(LIB)

clean cleanup:;

install: all;

.PRECIOUS:    $(LIB)
```

```
#       @(#)/usr/src/cmd/make/make.tm      3.2
#       ucb/io/70.mk  makefile
VERSION = 70

LIB = lib2.$(VERSION).a
COMPOOL=

LIB2OBJS = \
        $(LIB)(mx1.o) \
        $(LIB)(mx2.o) \
        $(LIB)(bio.o) \
        $(LIB)(tty.o) \
        $(LIB)(malloc.o) \
        $(LIB)(pipe.o) \
        $(LIB)(dhdm.o) \
        $(LIB)(dh.o) \
        $(LIB)(dhfdm.o) \
        $(LIB)(dj.o) \
        $(LIB)(dn.o) \
        $(LIB)(ds40.o) \
        $(LIB)(dz.o) \
        $(LIB)(alarm.o) \
        $(LIB)(hf.o) \
        $(LIB)(hps.o) \
        $(LIB)(hpmap.o) \
        $(LIB)(hp45.o) \
        $(LIB)(hs.o) \
        $(LIB)(ht.o) \
        $(LIB)(jy.o) \
        $(LIB)(kl.o) \
        $(LIB)(lfh.o) \
        $(LIB)(lp.o) \
        $(LIB)(mem.o) \
        $(LIB)(nmpipe.o) \
        $(LIB)(rf.o) \
        $(LIB)(rk.o) \
        $(LIB)(rp.o) \
        $(LIB)(rx.o) \
        $(LIB)(sys.o) \
        $(LIB)(trans.o) \
        $(LIB)(ttdma.o) \
        $(LIB)(tec.o) \
        $(LIB)(tex.o) \
        $(LIB)(tm.o) \
        $(LIB)(vp.o) \
        $(LIB)(vs.o) \
        $(LIB)(vtlp.o) \
        $(LIB)(vt11.o) \
        $(LIB)(fakevtlp.o) \
        $(LIB)(vt61.o) \
        $(LIB)(vt100.o) \
        $(LIB)(vtmon.o) \
        $(LIB)(vtdbg.o) \
        $(LIB)(vtutil.o) \
```

```
                $(LIB)(vtast.o) \
                $(LIB)(partab.o) \
                $(LIB)(rh.o) \
                $(LIB)(devstart.o) \
                $(LIB)(dmc11.o) \
                $(LIB)(rop.o) \
                $(LIB)(ioctl.o) \
                $(LIB)(fakemx.o)

all:    $(LIB)
        @echo $(LIB) is now up-to-date.

$(LIB)::$(LIB2OBJS)

$(LIB2OBJS):    $(FRC)

FRC:
        rm  -f $(LIB)

clobber: cleanup
        -rm  -f $(LIB) *.o

clean cleanup:;

install: all;

.PRECIOUS:      $(LIB)

.s.a:
        $(AS) $(ASFLAGS)  -o $*.o $<
        ar rcv $@ $*.o
        rm $*.o




#       @(#)/usr/src/cmd/make/make.tm     3.2
#       ucb/head/sys/70.mk makefile

COMPOOL  =  /usr/include/sys

HEADERS  =  \
        $(COMPOOL)/buf.h \
        $(COMPOOL)/bufx.h \
        $(COMPOOL)/conf.h \
        $(COMPOOL)/confx.h \
        $(COMPOOL)/crtctl.h \
        $(COMPOOL)/dir.h \
        $(COMPOOL)/dm11.h \
        $(COMPOOL)/elog.h \
        $(COMPOOL)/file.h \
        $(COMPOOL)/filex.h \
        $(COMPOOL)/filsys.h \
        $(COMPOOL)/ino.h \
        $(COMPOOL)/inode.h \
        $(COMPOOL)/inodex.h \
        $(COMPOOL)/ioctl.h \
```

```
                    $(COMPOOL)/ipcomm.h \
                    $(COMPOOL)/ipcommx.h \
                    $(COMPOOL)/lfsh.h \
                    $(COMPOOL)/lock.h \
                    $(COMPOOL)/maus.h \
                    $(COMPOOL)/mx.h \
                    $(COMPOOL)/param.h \
                    $(COMPOOL)/proc.h \
                    $(COMPOOL)/procx.h \
                    $(COMPOOL)/reg.h \
                    $(COMPOOL)/seg.h \
                    $(COMPOOL)/sgtty.h \
                    $(COMPOOL)/sigdef.h \
                    $(COMPOOL)/sprof.h \
                    $(COMPOOL)/sprofx.h \
                    $(COMPOOL)/stat.h \
                    $(COMPOOL)/syserr.h \
                    $(COMPOOL)/sysmes.h \
                    $(COMPOOL)/sysmesx.h \
                    $(COMPOOL)/systm.h \
                    $(COMPOOL)/text.h \
                    $(COMPOOL)/textx.h \
                    $(COMPOOL)/timeb.h \
                    $(COMPOOL)/trans.h \
                    $(COMPOOL)/tty.h \
                    $(COMPOOL)/ttyx.h \
                    $(COMPOOL)/types.h \
                    $(COMPOOL)/user.h \
                    $(COMPOOL)/userx.h \
                    $(COMPOOL)/version.h \
                    $(COMPOOL)/votrax.h \
                    $(COMPOOL)/vt11.h \
                    $(COMPOOL)/vtmn.h

all:     $(FRC) $(HEADERS)
         @echo Headers are now up-to-date.

$(HEADERS): s.$$(@F)
         $(GET) -s -p $(GFLAGS) $? > xtemp
         move xtemp 444 src sys $@

FRC:
         rm -f $(HEADERS)

.PRECIOUS:    $(HEADERS)

.h~.h:
         get -s $<

.DEFAULT:
         cpmv $? 444 src sys $@
```

A Tutorial Introduction to ADB

J. F. Maranzano

S. R. Bourne

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

Debugging tools generally provide a wealth of information about the inner workings of programs. These tools have been available on UNIX† to allow users to examine "core" files that result from aborted programs. A new debugging program, ADB, provides enhanced capabilities to examine "core" and other program files in a variety of formats, run programs with embedded breakpoints and patch files.

ADB is an indispensable but complex tool for debugging crashed systems and/or programs. This document provides an introduction to ADB with examples of its use. It explains the various formatting options, techniques for debugging C programs, examples of printing file system information and patching.

May 5, 1977

A Tutorial Introduction to ADB

J. F. Maranzano

S. R. Bourne

Bell Laboratories
Murray Hill, New Jersey 07974

1. Introduction

ADB is a new debugging program that is available on UNIX. It provides capabilities to look at "core" files resulting from aborted programs, print output in a variety of formats, patch files, and run programs with embedded breakpoints. This document provides examples of the more useful features of ADB. The reader is expected to be familiar with the basic commands on UNIX† with the C language, and with References 1, 2 and 3.

2. A Quick Survey

2.1. Invocation

ADB is invoked as:

> **adb objfile corefile**

where *objfile* is an executable UNIX file and *corefile* is a core image file. Many times this will look like:

> **adb a.out core**

or more simply:

> **adb**

where the defaults are *a.out* and *core* respectively. The filename minus (−) means ignore this argument as in:

> **adb − core**

ADB has requests for examining locations in either file. The ? request examines the contents of *objfile*, the / request examines the *corefile*. The general form of these requests is:

> **address ? format**

or

> **address / format**

2.2. Current Address

ADB maintains a current address, called dot, similar in function to the current pointer in the UNIX editor. When an address is entered, the current address is set to that location, so that:

> **0126?i**

---

†UNIX is a Trademark of Bell Laboratories.

sets dot to octal 126 and prints the instruction at that address. The request:

> ., 10/d

prints 10 decimal numbers starting at dot. Dot ends up referring to the address of the last item printed. When used with the ? or / requests, the current address can be advanced by typing newline; it can be decremented by typing ˆ.

Addresses are represented by expressions. Expressions are made up from decimal, octal, and hexadecimal integers, and symbols from the program under test. These may be combined with the operators +, −, *, % (integer division), & (bitwise and), | (bitwise inclusive or), # (round up to the next multiple), and ˜ (not). (All arithmetic within ADB is 32 bits.) When typing a symbolic address for a C program, the user can type *name* or *_name;* ADB will recognize both forms.

2.3. Formats

To print data, a user specifies a collection of letters and characters that describe the format of the printout. Formats are "remembered" in the sense that typing a request without one will cause the new printout to appear in the previous format. The following are the most commonly used format letters.

b	**one byte in octal**
c	**one byte as a character**
o	**one word in octal**
d	**one word in decimal**
f	**two words in floating point**
i	**PDP 11 instruction**
s	**a null terminated character string**
a	**the value of dot**
u	**one word as unsigned integer**
n	**print a newline**
r	**print a blank space**
ˆ	**backup dot**

(Format letters are also available for "long" values, for example, 'D' for long decimal, and 'F' for double floating point.) For other formats see the ADB manual.

2.4. General Request Meanings

The general form of a request is:

> **address,count command modifier**

which sets 'dot' to *address* and executes the command *count* times.

The following table illustrates some general ADB command meanings:

Command	Meaning
?	**Print contents from *a.out* file**
/	**Print contents from *core* file**
=	**Print value of "dot"**
:	**Breakpoint control**
$	**Miscellaneous requests**
;	**Request separator**
!	**Escape to shell**

ADB catches signals, so a user cannot use a quit signal to exit from ADB. The request $q or $Q (or cntl-D) must be used to exit from ADB.

3. Debugging C Programs

3.1. Debugging A Core Image

Consider the C program in Figure 1. The program is used to illustrate a common error made by C programmers. The object of the program is to change the lower case "t" to upper case in the string pointed to by *charp* and then write the character string to the file indicated by argument 1. The bug shown is that the character "T" is stored in the pointer *charp* instead of the string pointed to by *charp*. Executing the program produces a core file because of an out of bounds memory reference.

ADB is invoked by:

adb a.out core

The first debugging request:

$c

is used to give a C backtrace through the subroutines called. As shown in Figure 2 only one function (*main*) was called and the arguments *argc* and *argv* have octal values 02 and 0177762 respectively. Both of these values look reasonable; 02 = two arguments, 0177762 = address on stack of parameter vector.

The next request:

$C

is used to give a C backtrace plus an interpretation of all the local variables in each function and their values in octal. The value of the variable *cc* looks incorrect since *cc* was declared as a character.

The next request:

$r

prints out the registers including the program counter and an interpretation of the instruction at that location.

The request:

$e

prints out the values of all external variables.

A map exists for each file handled by ADB. The map for the *a.out* file is referenced by ? whereas the map for *core* file is referenced by /. Furthermore, a good rule of thumb is to use ? for instructions and / for data when looking at programs. To print out information about the maps type:

$m

This produces a report of the contents of the maps. More about these maps later.

In our example, it is useful to see the contents of the string pointed to by *charp*. This is done by:

*charp/s

which says use *charp* as a pointer in the *core* file and print the information as a character string. This printout clearly shows that the character buffer was incorrectly overwritten and helps identify the error. Printing the locations around *charp* shows that the buffer is unchanged but that the pointer is destroyed. Using ADB similarly, we could print information about the arguments to a function. The request:

main.argc/d

prints the decimal *core* image value of the argument *argc* in the function *main*.

The request:

*main.argv,3/o

prints the octal values of the three consecutive cells pointed to by *argv* in the function *main*. Note that these values are the addresses of the arguments to main. Therefore:

0177770/s

prints the ASCII value of the first argument. Another way to print this value would have been

*"/s

The " means ditto which remembers the last address typed, in this case *main.argc* ; the * instructs ADB to use the address field of the *core* file as a pointer.

The request:

.=o

prints the current address (not its contents) in octal which has been set to the address of the first argument. The current address, dot, is used by ADB to "remember" its current location. It allows the user to reference locations relative to the current address, for example:

.−10/d

3.2. Multiple Functions

Consider the C program illustrated in Figure 3. This program calls functions *f, g,* and *h* until the stack is exhausted and a core image is produced.

Again you can enter the debugger via:

adb

which assumes the names *a.out* and *core* for the executable file and core image file respectively. The request:

\$c

will fill a page of backtrace references to *f, g,* and *h*. Figure 4 shows an abbreviated list (typing *DEL* will terminate the output and bring you back to ADB request level).

The request:

,5\$C

prints the five most recent activations.

Notice that each function (*f,g,h*) has a counter of the number of times it was called.

The request:

fcnt/d

prints the decimal value of the counter for the function *f*. Similarly *gcnt* and *hcnt* could be printed. To print the value of an automatic variable, for example the decimal value of *x* in the last call of the function *h*, type:

h.x/d

It is currently not possible in the exported version to print stack frames other than the most recent activation of a function. Therefore, a user can print everything with \$C or the occurrence of a variable in the most recent call of a function. It is possible with the \$C request, however, to print the stack frame starting at some address as **address\$C**.

3.3. Setting Breakpoints

· Consider the C program in Figure 5. This program, which changes tabs into blanks, is adapted from *Software Tools* by Kernighan and Plauger, pp. 18-27.

We will run this program under the control of ADB (see Figure 6a) by:

adb a.out −

Breakpoints are set in the program as:

address:b [request]

The requests:

settab+4:b
fopen+4:b
getc+4:b
tabpos+4:b

set breakpoints at the start of these functions. C does not generate statement labels. Therefore it is currently not possible to plant breakpoints at locations other than function entry points without a knowledge of the code generated by the C compiler. The above addresses are entered as **symbol+4** so that they will appear in any C backtrace since the first instruction of each function is a call to the C save routine (*csv*). Note that some of the functions are from the C library.

To print the location of breakpoints one types:

$b

The display indicates a *count* field. A breakpoint is bypassed *count* −*1* times before causing a stop. The *command* field indicates the ADB requests to be executed each time the breakpoint is encountered. In our example no *command* fields are present.

By displaying the original instructions at the function *settab* we see that the breakpoint is set after the jsr to the C save routine. We can display the instructions using the ADB request:

settab,5?ia

This request displays five instructions starting at *settab* with the addresses of each location displayed. Another variation is:

settab,5?i

which displays the instructions with only the starting address.

Notice that we accessed the addresses from the *a.out* file with the ? command. In general when asking for a printout of multiple items, ADB will advance the current address the number of bytes necessary to satisfy the request; in the above example five instructions were displayed and the current address was advanced 18 (decimal) bytes.

To run the program one simply types:

:r

To delete a breakpoint, for instance the entry to the function *settab*, one types:

settab+4:d

To continue execution of the program from the breakpoint type:

:c

Once the program has stopped (in this case at the breakpoint for *fopen*), ADB requests can be used to display the contents of memory. For example:

$C

to display a stack trace, or:

tabs,3/8o

to print three lines of 8 locations each from the array called *tabs*. By this time (at location *fopen)* in the C program, *settab* has been called and should have set a one in every eighth location of *tabs*.

3.4. Advanced Breakpoint Usage

We continue execution of the program with:

:c

See Figure 6b. *Getc* is called three times and the contents of the variable *c* in the function *main* are displayed each time. The single character on the left hand edge is the output from the C program. On the third occurrence of *getc* the program stops. We can look at the full buffer of characters by typing:

ibuf+6/20c

When we continue the program with:

:c

we hit our first breakpoint at *tabpos* since there is a tab following the "This" word of the data.

Several breakpoints of *tabpos* will occur until the program has changed the tab into equivalent blanks. Since we feel that *tabpos* is working, we can remove the breakpoint at that location by:

tabpos+4:d

If the program is continued with:

:c

it resumes normal execution after ADB prints the message

a.out:running

The UNIX quit and interrupt signals act on ADB itself rather than on the program being debugged. If such a signal occurs then the program being debugged is stopped and control is returned to ADB. The signal is saved by ADB and is passed on to the test program if:

:c

is typed. This can be useful when testing interrupt handling routines. The signal is not passed on to the test program if:

:c 0

is typed.

Now let us reset the breakpoint at *settab* and display the instructions located there when we reach the breakpoint. This is accomplished by:

settab+4:b settab,5?ia *

It is also possible to execute the ADB requests for each occurrence of the breakpoint but only

---

* Owing to a bug in early versions of ADB (including the version distributed in Generic 3 UNIX) these statements must be written as:

```
        settab+4:b        settab,5?ia;0
        getc+4,3:b        main.c?C;0
        settab+4:b        settab,5?ia; ptab/o;0
```

Note that ;0 will set dot to zero and stop at the breakpoint.

stop after the third occurrence by typing:

getc+4,3:b main.c?C *

This request will print the local variable *c* in the function *main* at each occurrence of the breakpoint. The semicolon is used to separate multiple ADB requests on a single line.

Warning: setting a breakpoint causes the value of dot to be changed; executing the program under ADB does not change dot. Therefore:

settab+4:b .,5?ia
fopen+4:b

will print the last thing dot was set to (in the example *fopen+4*) *not* the current location (*settab+4*) at which the program is executing.

A breakpoint can be overwritten without first deleting the old breakpoint. For example:

settab+4:b settab,5?ia; ptab/o *

could be entered after typing the above requests.

Now the display of breakpoints:

$b

shows the above request for the *settab* breakpoint. When the breakpoint at *settab* is encountered the ADB requests are executed. Note that the location at *settab+4* has been changed to plant the breakpoint; all the other locations match their original value.

Using the functions, *f, g* and *h* shown in Figure 3, we can follow the execution of each function by planting non-stopping breakpoints. We call ADB with the executable program of Figure 3 as follows:

adb ex3 −

Suppose we enter the following breakpoints:

```
h+4:b     hcnt/d; h.hi/; h.hr/
g+4:b     gcnt/d; g.gi/; g.gr/
f+4:b     fcnt/d; f.fi/; f.fr/
:r
```

Each request line indicates that the variables are printed in decimal (by the specification **d**). Since the format is not changed, the **d** can be left off all but the first request.

The output in Figure 7 illustrates two points. First, the ADB requests in the breakpoint line are not examined until the program under test is run. That means any errors in those ADB requests is not detected until run time. At the location of the error ADB stops running the program.

The second point is the way ADB handles register variables. ADB uses the symbol table to address variables. Register variables, like *f.fr* above, have pointers to uninitialized places on the stack. Therefore the message "symbol not found".

Another way of getting at the data in this example is to print the variables used in the call as:

```
f+4:b     fcnt/d; f.a/; f.b/; f.fi/
g+4:b     gcnt/d; g.p/; g.q/; g.gi/
:c
```

The operator / was used instead of ? to read values from the *core* file. The output for each function, as shown in Figure 7, has the same format. For the function *f*, for example, it shows the name and value of the *external* variable *fcnt*. It also shows the address on the stack and value of the variables *a, b* and *fi*.

Notice that the addresses on the stack will continue to decrease until no address space is left for program execution at which time (after many pages of output) the program under test aborts. A display with names would be produced by requests like the following:

f+4:b   fcnt/d; f.a/"a="d; f.b/"b="d; f.fi/"fi="d

In this format the quoted string is printed literally and the **d** produces a decimal display of the variables. The results are shown in Figure 7.

3.5. Other Breakpoint Facilities

● Arguments and change of standard input and output are passed to a program as:

:r arg1 arg2 ... <infile >outfile

This request kills any existing program under test and starts the *a.out* afresh.

● The program being debugged can be single stepped by:

:s

If necessary, this request will start up the program being debugged and stop after executing the first instruction.

● ADB allows a program to be entered at a specific address by typing:

address:r

● The count field can be used to skip the first *n* breakpoints as:

,n:r

The request:

,n:c

may also be used for skipping the first *n* breakpoints when continuing a program.

● A program can be continued at an address different from the breakpoint by:

address:c

● The program being debugged runs as a separate process and can be killed by:

:k

4. Maps

UNIX supports several executable file formats. These are used to tell the loader how to load the program file. File type 407 is the most common and is generated by a C compiler invocation such as **cc pgm.c.** A 410 file is produced by a C compiler command of the form **cc -n pgm.c,** whereas a 411 file is produced by **cc -i pgm.c.** ADB interprets these different file formats and provides access to the different segments through a set of maps (see Figure 8). To print the maps type:

$m

In 407 files, both text (instructions) and data are intermixed. This makes it impossible for ADB to differentiate data from instructions and some of the printed symbolic addresses look incorrect; for example, printing data addresses as offsets from routines.

In 410 files (shared text), the instructions are separated from data and **?*** accesses the data part of the *a.out* file. The **?*** request tells ADB to use the second part of the map in the *a.out* file. Accessing data in the *core* file shows the data after it was modified by the execution

of the program. Notice also that the data segment may have grown during program execution.

· In 411 files (separated I & D space), the instructions and data are also separated. However, in this case, since data is mapped through a separate set of segmentation registers, the base of the data segment is also relative to address zero. In this case since the addresses overlap it is necessary to use the ?* operator to access the data space of the *a.out* file. In both 410 and 411 files the corresponding core file does not contain the program text.

Figure 9 shows the display of three maps for the same program linked as a 407, 410, 411 respectively. The b, e, and f fields are used by ADB to map addresses into file addresses. The "f1" field is the length of the header at the beginning of the file (020 bytes for an *a.out* file and 02000 bytes for a *core* file). The "f2" field is the displacement from the beginning of the file to the data. For a 407 file with mixed text and data this is the same as the length of the header; for 410 and 411 files this is the length of the header plus the size of the text portion.

The "b" and "e" fields are the starting and ending locations for a segment. Given an address, A, the location in the file (either *a.out* or *core*) is calculated as:

$$b1 \leqslant A \leqslant e1 \Rightarrow \text{file address} = (A-b1)+f1$$
$$b2 \leqslant A \leqslant e2 \Rightarrow \text{file address} = (A-b2)+f2$$

A user can access locations by using the ADB defined variables. The $v request prints the variables initialized by ADB:

b	base address of data segment
d	length of the data segment
s	length of the stack
t	length of the text
m	execution type (407,410,411)

In Figure 9 those variables not present are zero. Use can be made of these variables by expressions such as:

    <b

in the address field. Similarly the value of the variable can be changed by an assignment request such as:

    02000>b

that sets b to octal 2000. These variables are useful to know if the file under examination is an executable or *core* image file.

ADB reads the header of the *core* image file to find the values for these variables. If the second file specified does not seem to be a *core* file, or if it is missing then the header of the executable file is used instead.

5. Advanced Usage

It is possible with ADB to combine formatting requests to provide elaborate displays. Below are several examples.

5.1. Formatted dump

The line:

    <b,−1/4o4^8Cn

prints 4 octal words followed by their ASCII interpretation from the data space of the core image file. Broken down, the various request pieces mean:

    <b   The base address of the data segment.

<b,−1   Print from the base address to the end of file. A negative count is used here and elsewhere to loop indefinitely or until some error condition (like end of file) is detected.

The format **4o4^8Cn** is broken down as follows:

4o   Print 4 octal locations.

4^   Backup the current address 4 locations (to the original start of the field).

8C   Print 8 consecutive characters using an escape convention; each character in the range 0 to 037 is printed as @ followed by the corresponding character in the range 0140 to 0177. An @ is printed as @@.

n   Print a newline.

The request:

<b,<d/4o4^8Cn

could have been used instead to allow the printing to stop at the end of the data segment (<d provides the data segment size in bytes).

The formatting requests can be combined with ADB's ability to read in a script to produce a core image dump script. ADB is invoked as:

adb a.out core < dump

to read in a script file, *dump*, of requests. An example of such a script is:

```
120$w
4095$s
$v
=3n
$m
=3n"C Stack Backtrace"
$C
=3n"C External Variables"
$e
=3n"Registers"
$r
0$s
=3n"Data Segment"
<b,−1/8ona
```

The request **120$w** sets the width of the output to 120 characters (normally, the width is 80 characters). ADB attempts to print addresses as:

symbol + offset

The request **4095$s** increases the maximum permissible offset to the nearest symbolic address from 255 (default) to 4095. The request = can be used to print literal strings. Thus, headings are provided in this *dump* program with requests of the form:

=3n"C Stack Backtrace"

that spaces three lines and prints the literal string. The request **$v** prints all non-zero ADB variables (see Figure 8). The request **0$s** sets the maximum offset for symbol matches to zero

thus suppressing the printing of symbolic labels in favor of octal values. Note that this is only done for the printing of the data segment. The request:

<b,-1/8ona

prints a dump from the base of the data segment to the end of file with an octal address field and eight octal numbers per line.

Figure 11 shows the results of some formatting requests on the C program of Figure 10.

5.2. Directory Dump

As another illustration (Figure 12) consider a set of requests to dump the contents of a directory (which is made up of an integer *inumber* followed by a 14 character name):

```
. adb dir -
=n8t"Inum"8t"Name"
0,-1? u8t14cn
```

In this example, the u prints the *inumber* as an unsigned decimal integer, the 8t means that ADB will space to the next multiple of 8 on the output line, and the 14c prints the 14 character file name.

5.3. Ilist Dump

. Similarly the contents of the *ilist* of a file system, (e.g. /dev/src, on UNIX systems distributed by the UNIX Support Group; see UNIX Programmer's Manual Section V) could be dumped with the following set of requests:

```
adb /dev/src -
02000>b
?m <b
<b,-1?"flags"8ton"links,uid,gid"8t3bn",size"8tbrdn"addr"8t8un"times"8t2Y2na
```

In this example the value of the base for the map was changed to 02000 (by saying **?m<b**) since that is the start of an *ilist* within a file system. An artifice (**brd** above) was used to print the 24 bit size field as a byte, a space, and a decimal integer. The last access time and last modify time are printed with the 2Y operator. Figure 12 shows portions of these requests as applied to a directory and file system.

5.4. Converting values

ADB may be used to convert values from one representation to another. For example:

072 = odx

will print

072   58   #3a

which is the octal, decimal and hexadecimal representations of 072 (octal). The format is remembered so that typing subsequent numbers will print them in the given formats. Character values may be converted similarly, for example:

'a' = co

prints

a   0141

It may also be used to evaluate expressions but be warned that all binary operators have the same precedence which is lower than that for unary operators.

6. Patching

Patching files with ADB is accomplished with the *write*, w or W, request (which is not like the *ed* editor write command). This is often used in conjunction with the *locate*, l or L request. In general, the request syntax for l and w are similar as follows:

?l value

The request l is used to match on two bytes, L is used for four bytes. The request w is used to write two bytes, whereas W writes four bytes. The **value** field in either *locate* or *write* requests is an expression. Therefore, decimal and octal numbers, or character strings are supported

In order to modify a file, ADB must be called as: '

adb −w file1 file2

When called with this option, *file1* and *file2* are created if necessary and opened for both reading and writing.

For example, consider the C program shown in Figure 10. We can change the word **"This"** to **"The "** in the executable file for this program, *ex7*, by using the following requests:

```
adb −w ex7 −
?l 'Th'
?W 'The '
```

The request ?l starts at dot and stops at the first match of "Th" having set dot to the address of the location found. Note the use of ? to write to the *a.out* file. The form ?* would have been used for a 411 file.

More frequently the request will be typed as:

?l 'Th'; ?s

and locates the first occurrence of "Th" and print the entire string. Execution of this ADB request will set dot to the address of the "Th" characters.

As another example of the utility of the patching facility, consider a C program that has an internal logic flag. The flag could be set by the user through ADB and the program run. For example:

```
adb a.out −
:s arg1 arg2
flag/w 1
:c
```

The :s request is normally used to single step through a process or start a process in single step mode. In this case it starts *a.out* as a subprocess with arguments **arg1** and **arg2**. If there is a subprocess running ADB writes to it rather than to the file so the w request causes *flag* to be changed in the memory of the subprocess.

7. Anomalies

Below is a list of some strange things that users should be aware of.

1. Function calls and arguments are put on the stack by the C save routine. Putting breakpoints at the entry point to routines means that the function appears not to have been called when the breakpoint occurs.

2. When printing addresses, ADB uses either text or data symbols from the *a.out* file. This sometimes causes unexpected symbol names to be printed with data (e.g. *savr5+022*). This does not happen if ? is used for text (instructions) and / for data.

3.   ADB cannot handle C register variables in the most recently activated function.

8. Acknowledgements

The authors are grateful for the thoughtful comments on how to organize this document from R. B. Brandt, E. N. Pinson and B. A. Tague. D. M. Ritchie made the system changes necessary to accommodate tracing within ADB. He also participated in discussions during the writing of ADB. His earlier work with DB and CDB led to many of the features found in ADB.

9. References

1.   D. M. Ritchie and K. Thompson, "The UNIX Time-Sharing System," CACM, July, 1974.
2.   B. W. Kernighan and D. M. Ritchie, *The C Programming Language,* Prentice-Hall, 1978.
3.   K. Thompson and D. M. Ritchie, UNIX Programmer's Manual - 7th Edition, 1978.
4.   B. W. Kernighan and P. J. Plauger, *Software Tools,* Addison-Wesley, 1976.

Figure 1: C program with pointer bug

```
struct buf {
        int fildes;
        int nleft;
        char *nextp;
        char buff[512];
        }bb;
struct buf *obuf;

char *charp "this is a sentence.";

main(argc,argv)
int argc;
char **argv;
{
        char    cc;

        if(argc < 2) {
                printf("Input file missing\n");
                exit(8);
        }

        if((fcreat(argv[1],obuf)) < 0){
                printf("%s : not found\n", argv[1]);
                exit(8);
        }
        charp = 'T';
printf("debug 1 %s\n",charp);
        while(cc= *charp++)
                putc(cc,obuf);
        fflush(obuf);
}
```

Figure 2: ADB output for C program of Figure 1

```
adb a.out core
$c
~main(02,0177762)
$C
~main(02,0177762)
          argc:       02
          argv:       0177762
          cc:         02124
$r
ps        0170010
pc        0204      ~main+0152
sp        0177740
r5        0177752
r4        01
r3        0
r2        0
r1        0
r0        0124
~main+0152:    mov     _obuf,(sp)
$e
savr5:    0
_obuf:    0
_charp:   0124
_errno:   0
_fout:    0
$m
text map   `ex1'
b1 = 0              e1  = 02360          f1 = 020
b2 = 0              e2  = 02360          f2 = 020
data map   `core1'
b1 = 0              e1  = 03500          f1 = 02000
b2 = 0175400        e2  = 0200000        f2 = 05500
*charp/s
0124:          TTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTTLx        Nh@x&_
-

charp/s
_charp:        T

_charp+02:     this is a sentence.

_charp+026:    Input file missing
main.argc/d
0177756:       2
*main.argv/3o
0177762:       0177770 0177776 0177777
0177770/s
0177770:       a.out
*main.argv/3o
0177762:       0177770 0177776 0177777
*"/s
0177770:       a.out
 .=o
               0177770
 .-10/d
0177756:       2
$q
```

Figure 3: Multiple function C program for stack trace illustration

```
int      fcnt,gcnt,hcnt;
h(x,y)
{
         int hi; register int hr;
         hi = x+1;
         hr = x−y+1;
         hcnt++ ;
         hj:
         f(hr,hi);
}

g(p,q)
{
         int gi; register int gr;
         gi = q−p;
         gr = q−p+1;
         gcnt++ ;
         gj:
         h(gr,gi);
}

f(a,b)
{
         int fi; register int fr;
         fi = a+2*b;
         fr = a+b;
         fcnt++ ;
         fj:
         g(fr,fi);
}

main()
{
         f(1,1);
}
```

Figure 4: ADB output for C program of Figure 3

```
adb
$c
‾h(04452,04451)
‾g(04453,011124)
‾f(02,04451)
‾h(04450,04447)
‾g(04451,011120)
‾f(02,04447)
‾h(04446,04445)
‾g(04447,011114)
‾f(02,04445)
‾h(04444,04443)
HIT DEL KEY
adb
,5$C
‾h(04452,04451)
          x:          04452
          y:          04451
          hi:         ?
‾g(04453,011124)
          p:          04453
          q:          011124
          gi:         04451
          gr:         ?
‾f(02,04451)
          a:          02
          b:          04451
          fi:         011124
          fr:         04453
‾h(04450,04447)
          x:          04450
          y:          04447
          hi:         04451
          hr:         02
‾g(04451,011120)
          p:          04451
          q:          011120
          gi:         04447
          gr:         04450
fcnt/d
_fcnt:       1173
..gcnt/d
_gcnt:       1173
hcnt/d
_hcnt:       1172
h.x/d
022004:      2346
$q
```

Figure 5: C program to decode tabs

```
#define MAXLINE      80
#define YES          1
#define NO           0
#define TABSP        8

char    input[] "data";  .
char    ibuf[518];
int     tabs[MAXLINE];

main()
{
        int col, *ptab;
        char c;

        ptab = tabs;
        settab(ptab);       /*Set initial tab stops */
        col = 1;
        if(fopen(input,ibuf) < 0) {
                printf("%s : not found\n",input);
                exit(8);
        }
        while((c = getc(ibuf)) != -1) {
                switch(c) {
                        case '\t': /* TAB */
                                while(tabpos(col) != YES) {
                                        putchar(' ');       /* put BLANK */
                                        col++ ;
                                }
                                break;
                        case '\n':/*NEWLINE */
                                putchar('\n');
                                col = 1;
                                break;
                        default:
                                putchar(c);
                                col++ ;
                }
        }
}

/* Tabpos return YES if col is a tab stop */
tabpos(col)
int col;
{
        if(col > MAXLINE)
                return(YES);
        else
                return(tabs[col]);
}

/* Settab - Set initial tab stops */
settab(tabp)
int *tabp;
{
        int i;
        for(i = 0; i<= MAXLINE; i++)
                (i%TABSP) ? (tabs[i] = NO) : (tabs[i] = YES);
}
```

Figure 6a: ADB output for C program of Figure 5

```
adb a.out −
settab+4:b
fopen+4:b
getc+4:b
tabpos+4:b
$b
breakpoints
count   bkpt            command
1       ~tabpos+04
1       _getc+04
1       _fopen+04
1       ~settab+04
settab,5?ia
~settab:         jsr     r5,csv
~settab+04:      tst     −(sp)
~settab+06:      clr     0177770(r5)
~settab+012:     cmp     $0120,0177770(r5)
~settab+020:     blt     ~settab+076
~settab+022:
settab,5?i
~settab:         jsr     r5,csv
                 tst     −(sp)
                 clr     0177770(r5)
                 cmp     $0120,0177770(r5)
                 blt     ~settab+076
:r
a.out: running
breakpoint      ~settab+04:      tst     −(sp)
settab+4:d
:c
a.out: running
breakpoint      _fopen+04:      mov     04(r5),nulstr+012
$C
_fopen(02302,02472)
~main(01,0177770)
        col:     01
        c:       0
        ptab:    03500
tabs,3/8o
03500:          01      0   0   0   0   0   0   0
                01      0   0   0   0   0   0   0
                01      0   0   0   0   0   0   0
```
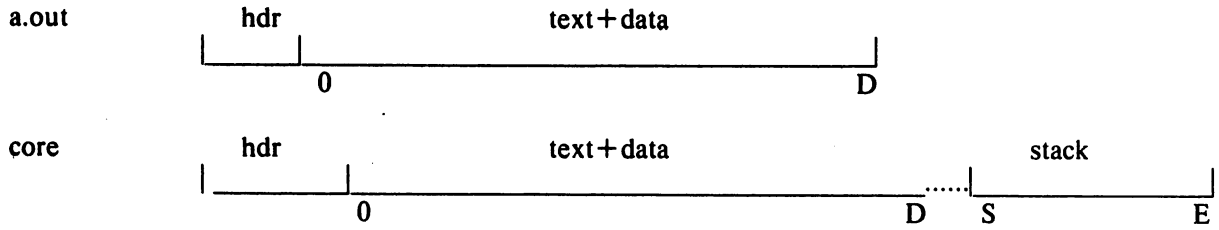
Figure 6b: ADB output for C program of Figure 5

```
:c
a.out: running
breakpoint        _getc+04:        mov     04(r5),r1
ibuf+6/20c
   _cleanu+0202:                   This     is       a test    of
:c
a.out: running
breakpoint        ~tabpos+04:      cmp     $0120,04(r5)
tabpos+4:d
settab+4:b  settab,5?ia
settab+4:b  settab,5?ia;  0
getc+4,3:b  main.c?C;  0
settab+4:b  settab,5?ia;  ptab/o;  0
$b
breakpoints
count    bkpt              command
1        ~tabpos+04
3        _getc+04          main.c?C;0
1        _fopen+04
1        ~settab+04        settab,5?ia;ptab?o;0
~settab:          jsr     r5,csv
~settab+04:       bpt
~settab+06:       clr     0177770(r5)
~settab+012:      cmp     $0120,0177770(r5)
~settab+020:      blt     ~settab+076
~settab+022:
0177766:          0177770
0177744:          @`
T0177744:         T
h0177744:         h
i0177744:         i
s0177744:         s
```

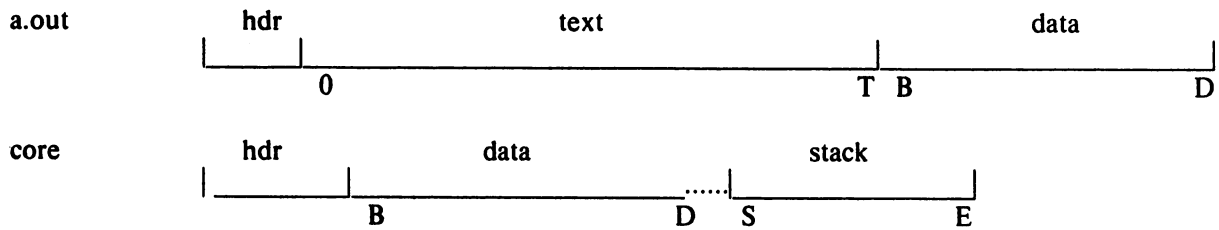**Figure 7: ADB output for C program with breakpoints**

```
adb ex3 −
h+4:b hcnt/d; h.hi/; h.hr/
g+4:b gcnt/d; g.gi/; g.gr/
f+4:b fcnt/d; f.fi/; f.fr/
:r
ex3: running
_fcnt:          0
0177732:        214
symbol not found
f+4:b fcnt/d; f.a/; f.b/; f.fi/
g+4:b gcnt/d; g.p/; g.q/; g.gi/
h+4:b hcnt/d; h.x/; h.y/; h.hi/
:c
ex3: running
_fcnt:          0
0177746:        1
0177750:        1
0177732:        214
_gcnt:          0
0177726:        2
0177730:        3
0177712:        214
_hcnt:          0
0177706:        2
0177710:        1
0177672:        214
_fcnt:          1
0177666:        2
0177670:        3
0177652:        214
_gcnt:          1
0177646:        5
0177650:        8
0177632:        214
HIT DEL
f+4:b fcnt/d; f.a/"a = "d; f.b/"b = "d; f.fi/"fi = "d
g+4:b gcnt/d; g.p/"p = "d; g.q/"q = "d; g.gi/"gi = "d
h+4:b hcnt/d; h.x/"x = "d; h.y/"h = "d; h.hi/"hi = "d
:r
ex3: running
_fcnt:          0
0177746:        a = 1
0177750:        b = 1
0177732:        fi = 214
_gcnt:          0
0177726:        p = 2
0177730:        q = 3
0177712:        gi = 214
_hcnt:          0
0177706:        x = 2
0177710:        y = 1
0177672:        hi = 214
_fcnt:          1
0177666:        a = 2
0177670:        b = 3
0177652:        fi = 214
HIT DEL
$q
```
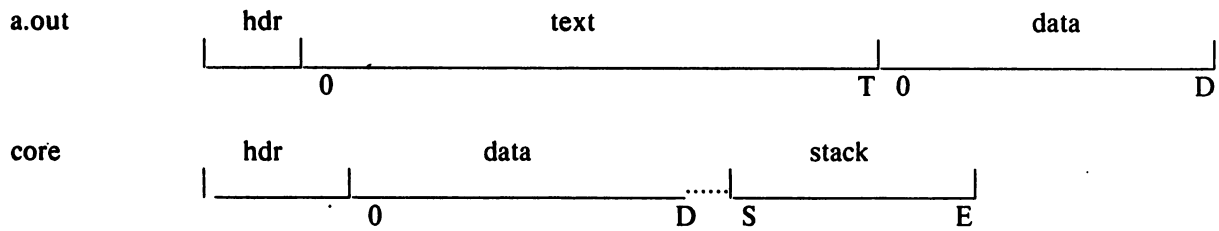
Figure 8: ADB address maps

407 files

a.out

| hdr | text + data |

0   D

core

| hdr | text + data | stack |

0   D   S   E

410 files (shared text)

a.out

| hdr | text | data |

0   T   B   D

core

| hdr | data | stack |

B   D   S   E

411 files (separated I and D space)

a.out

| hdr | text | data |

0   T   0   D

core

| hdr | data | stack |

0   D   S   E

The following *adb* variables are set.

		407	410	411
b	base of data	0	B	0
d	length of data	D	D−B	D
s	length of stack	S	S	S
t	length of text	0	T	T

Figure 9: ADB output for maps

```
adb map407 core407
$m
text map     `map407'
b1 = 0              e1    = 0256        f1 = 020
b2 = 0              e2    = 0256        f2 = 020
data map     `core407'
b1 = 0              e1    = 0300        f1 = 02000
b2 = 0175400        e2    = 0200000     f2 = 02300
$v
variables
d = 0300
m = 0407
s = 02400
$q



adb map410 core410
$m
text map     `map410'
b1 = 0              e1    = 0200        f1 = 020
b2 = 020000         e2    = 020116   f2 = 0220
data map     `core410'
b1 = 020000         e1    = 020200   f1 = 02000
b2 = 0175400        e2    = 0200000     f2 = 02200
$v
variables
b = 020000
d = 0200
m = 0410
s = 02400
t = 0200
$q



adb map411 core411
$m
text map     `map411'
b1 = 0              e1    = 0200        f1 = 020
b2 = 0              e2    = 0116        f2 = 0220
data map     `core411'
b1 = 0              e1    = 0200        f1 = 02000
b2 = 0175400        e2    = 0200000     f2 = 02200
$v
variables
d = 0200
m = 0411
s = 02400
t = 0200
$q
```

Figure 10: Simple C program for illustrating formatting and patching

```
char    str1[]    "This is a character string";
int     one       1;
int     number 456;
long    lnum      1234;
float   fpt       1.25;
char    str2[]    "This is the second character string";
main()
{
        one = 2;
}
```

Figure 11: ADB output illustrating fancy formats

```
adb map410 core410
<b,-1/8ona
020000:           0      064124    071551    064440    020163    020141    064143    071141

_str1+016: 061541     062564    020162    072163    064562    063556    0     02

_number:
_number: 0710 0      02322040240      0    064124    071551    064440

_str2+06: 020163     064164    020145    062563    067543    062156    061440    060550

_str2+026: 060562    072143    071145    071440    071164    067151    0147 0

savr5+02: 0     0     0     0     0     0     0     0

<b,20/4o4^8Cn
020000:           0      064124    071551    064440    @`@`This i
                  020163    020141    064143    071141    s a char
                  061541    062564    020162    072163    acter st
                  064562    063556    0     02    ring@`@`@b@`

_number: 0710 0      02322040240      H@a@`@`R@d @@
         0     064124    071551    064440    @`@`This i
                  020163    064164    020145    062563    s the se
                  067543    062156    061440    060550    cond cha
                  060562    072143    071145    071440    racter s
                  071164    067151    0147 0    tring@`@`@`
                  0     0     0     0     @`@`@`@`@`@`@`@`
                  0     0     0     0     @`@`@`@`@`@`@`@`
data address not found
<b,20/4o4^8t8cna
020000:           0      064124    071551    064440          This i
_str1+06: 020163     020141    064143    071141          s a char
_str1+016: 061541    062564    020162    072163          acter st
_str1+026: 064562    063556    0     02          ring
_number:
_number: 0710 0      02322040240          HR
_fpt+02: 0     064124    071551    064440          This i
_str2+06: 020163     064164    020145    062563          s the se
_str2+016: 067543    062156    061440    060550          cond cha
_str2+026: 060562    072143    071145    071440          racter s
_str2+036: 071164    067151    0147 0          tring
savr5+02: 0     0     0     0
savr5+012: 0     0     0     0
data address not found
<b,10/2b8t^2cn
020000:           0      0

_str1:    0124 0150      Th
          0151 0163      is
          040  0151      i
          0163 040       s
          0141 040       a
          0143 0150      ch
          0141 0162      ar
          0141 0143      ac
          0164 0145      te

$Q
```

Figure 12: Directory and inode dumps
adb dir −
=nt"Inode"t"Name"
0,−1?ut14cn

```
            Inode       Name
0:          652 .
            82   ..
            5971 cap.c
            5323 cap
            0    pp
```

```
adb /dev/src −
02000>b
?m<b
 new map      `/dev/src´
b1 = 02000         e1      = 0100000000   f1 = 0
b2 = 0             e2      = 0        f2 = 0
$v
variables
b = 02000
<b,−1?"flags"8ton"links,uid,gid"8t3bn"size"8tbrdn"addr"8t8un"times"8t2YZna
02000:            flags 073145
            links,uid,gid      0163 0164 0141
            size  0162 10356
            addr 28770  ·  8236 25956      27766      25455      8236 25956      25206
            times1976 Feb 5 08:34:56   1975 Dec 28 10:55:15


02040:            flags 024555
            links,uid,gid      012  0163 0164
            size  0162 25461
            addr 8308 30050       8294 25130      15216      26890      29806      10784
            times1976 Aug 17 12:16:51 1976 Aug 17 12:16:51


02100:            flags 05173
            links,uid,gid      011  0162 0145
            size  0147 29545
            addr 25972       8306 28265       8308 25642      15216      2314 25970
            times1977 Apr 2 08:58:01   1977 Feb 5 10:21:44
```
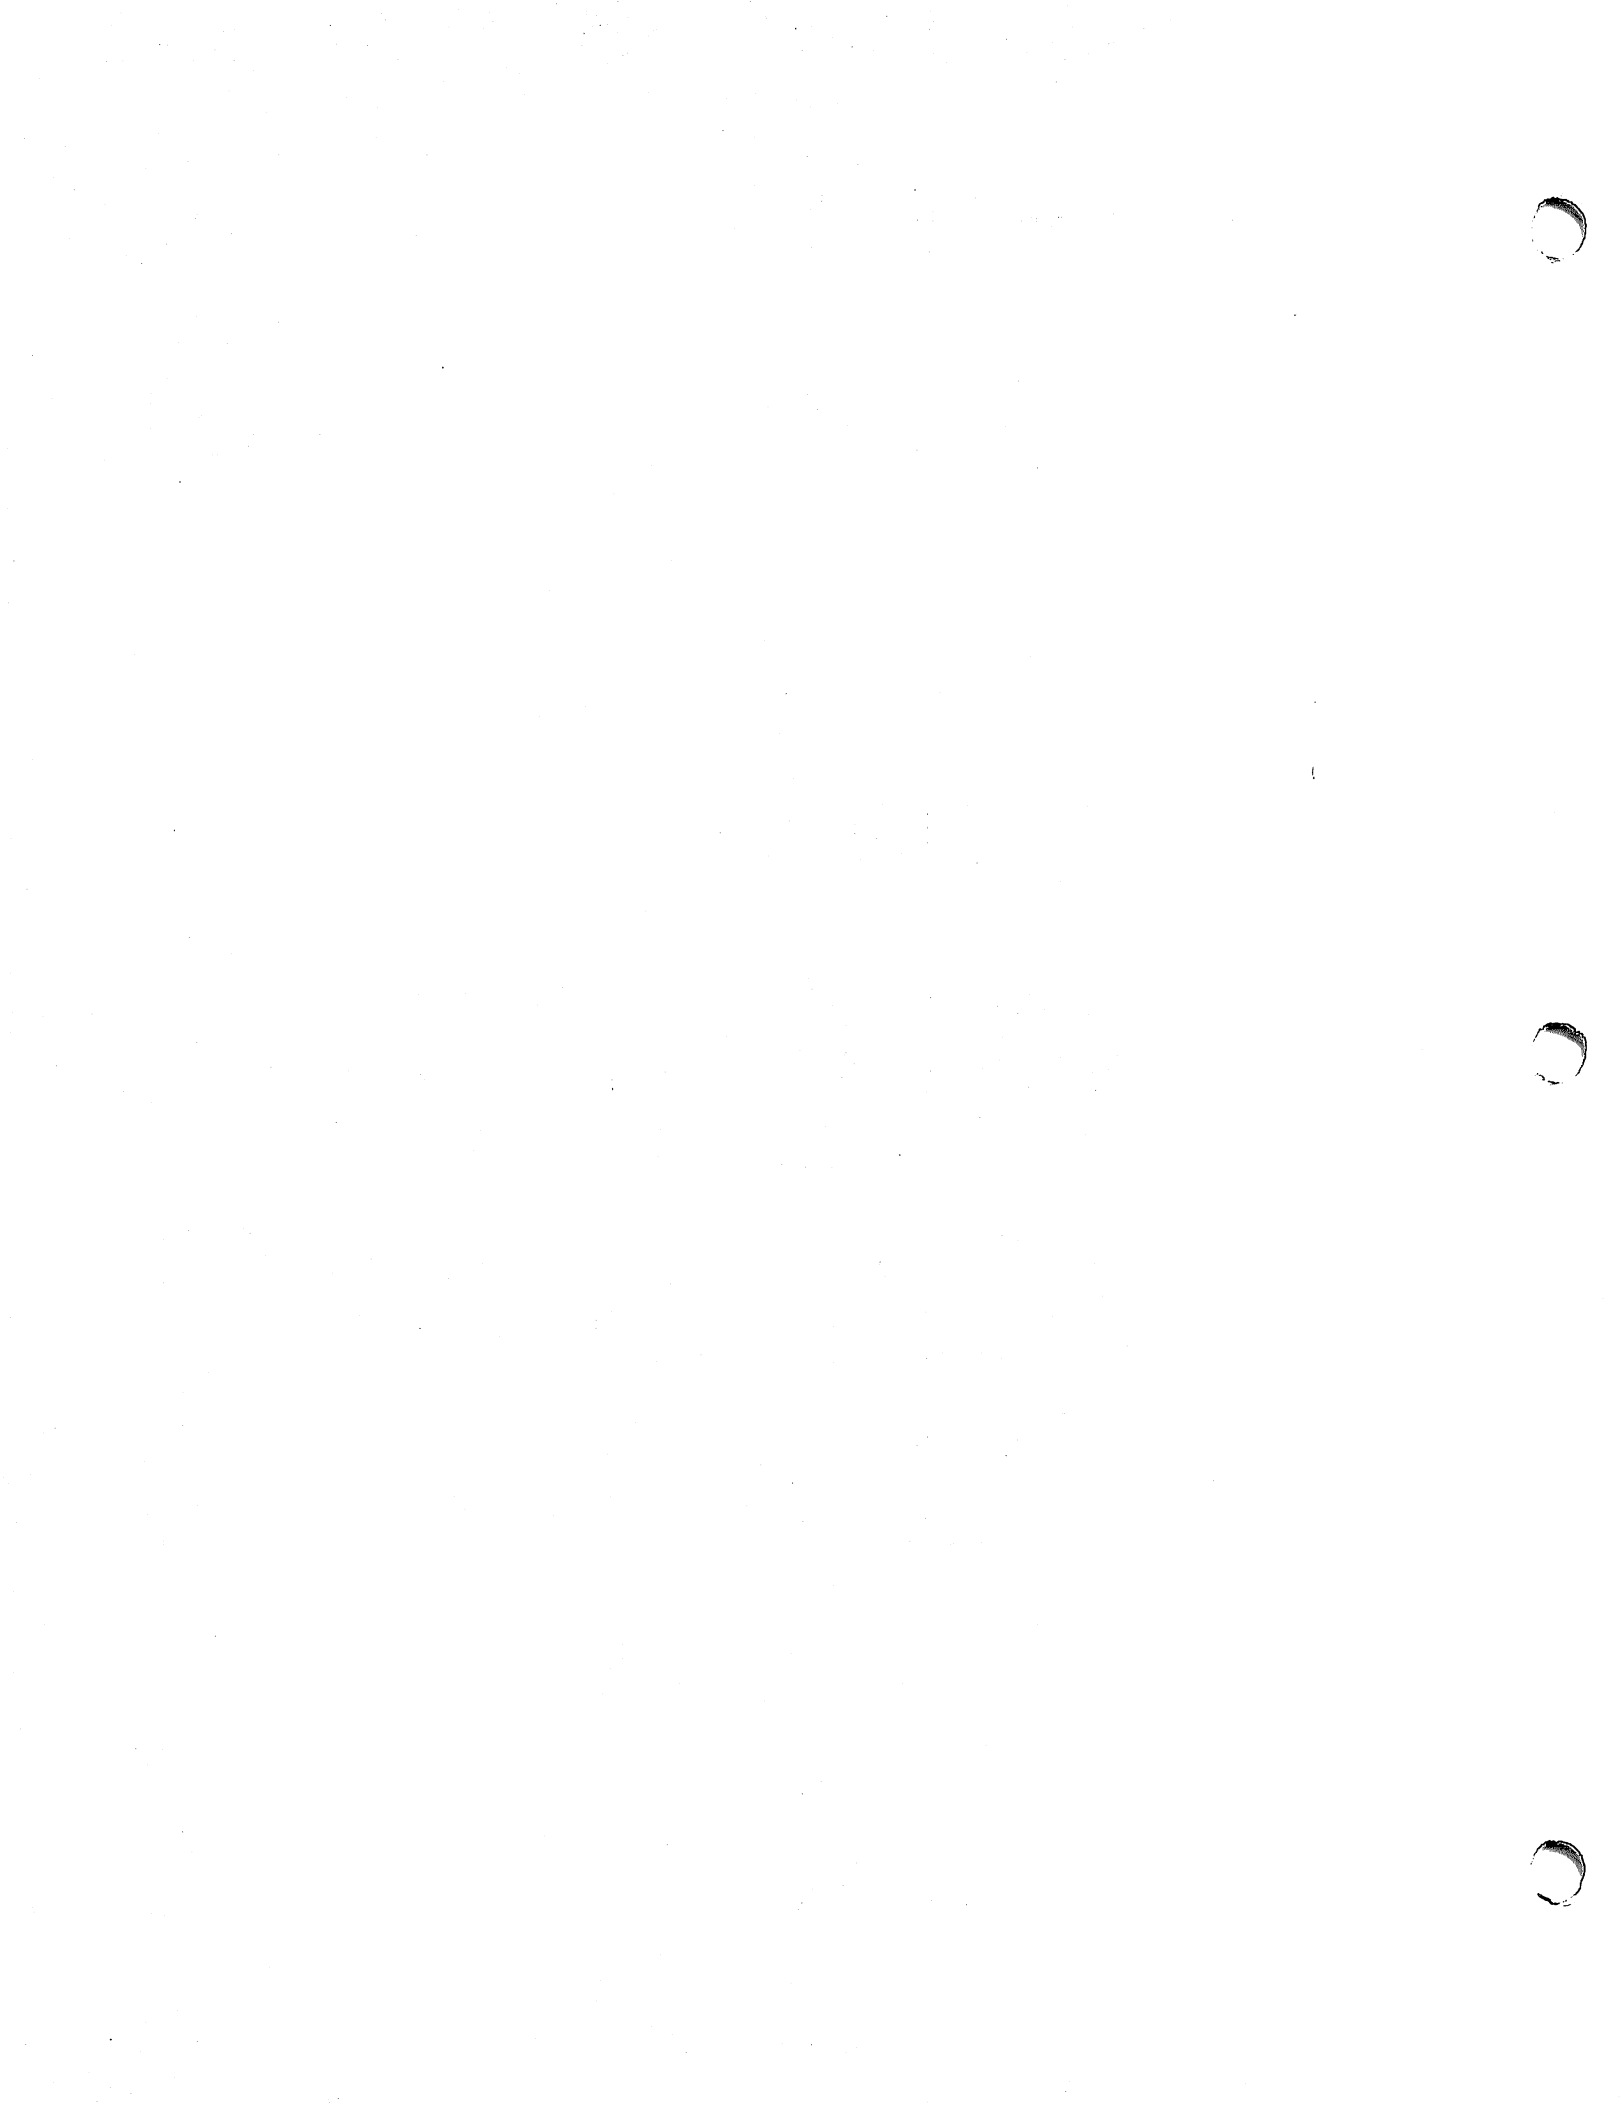
ADB Summary

Command Summary

a) formatted printing

? *format*   print from *a.out* file according to *format*

/ *format*   print from *core* file according to *format*

= *format*   print the value of *dot*

?w expr   write expression into *a.out* file

/w expr   write expression into *core* file

?l expr   locate expression in *a.out* file

b) breakpoint and program control

:b	set breakpoint at *dot*
:c	continue running program
:d	delete breakpoint
:k	kill the program being debugged
:r	run *a.out* file under ADB control
:s	single step

c) miscellaneous printing

$b	print current breakpoints
$c	C stack trace
$e	external variables
$f	floating registers
$m	print ADB segment maps
$q	exit from ADB
$r	general registers
$s	set offset for symbol match
$v	print ADB variables
$w	set output line width

d) calling the shell

!   call *shell* to read rest of line

e) assignment to variables

> *name*   assign dot to variable or register *name*

Format Summary

a	the value of dot
b	one byte in octal
c	one byte as a character
d	one word in decimal
f	two words in floating point
i	PDP 11 instruction
o	one word in octal
n	print a newline
r	print a blank space
s	a null terminated character string
*n*t	move to next *n* space tab
u	one word as unsigned integer
x	hexadecimal
Y	date
^	backup dot
"..."	print string

Expression Summary

a) expression components

decimal integer	e.g. 256
octal integer	e.g. 0277
hexadecimal	e.g. #ff
symbols	e.g. flag _main main.argc
variables	e.g. <b
registers	e.g. <pc <r0
(expression)	expression grouping

b) dyadic operators

+	add
−	subtract
*	multiply
%	integer division
&	bitwise and
|	bitwise or
#	round up to the next multiple

c) monadic operators

~	not
*	contents of location
−	integer negate
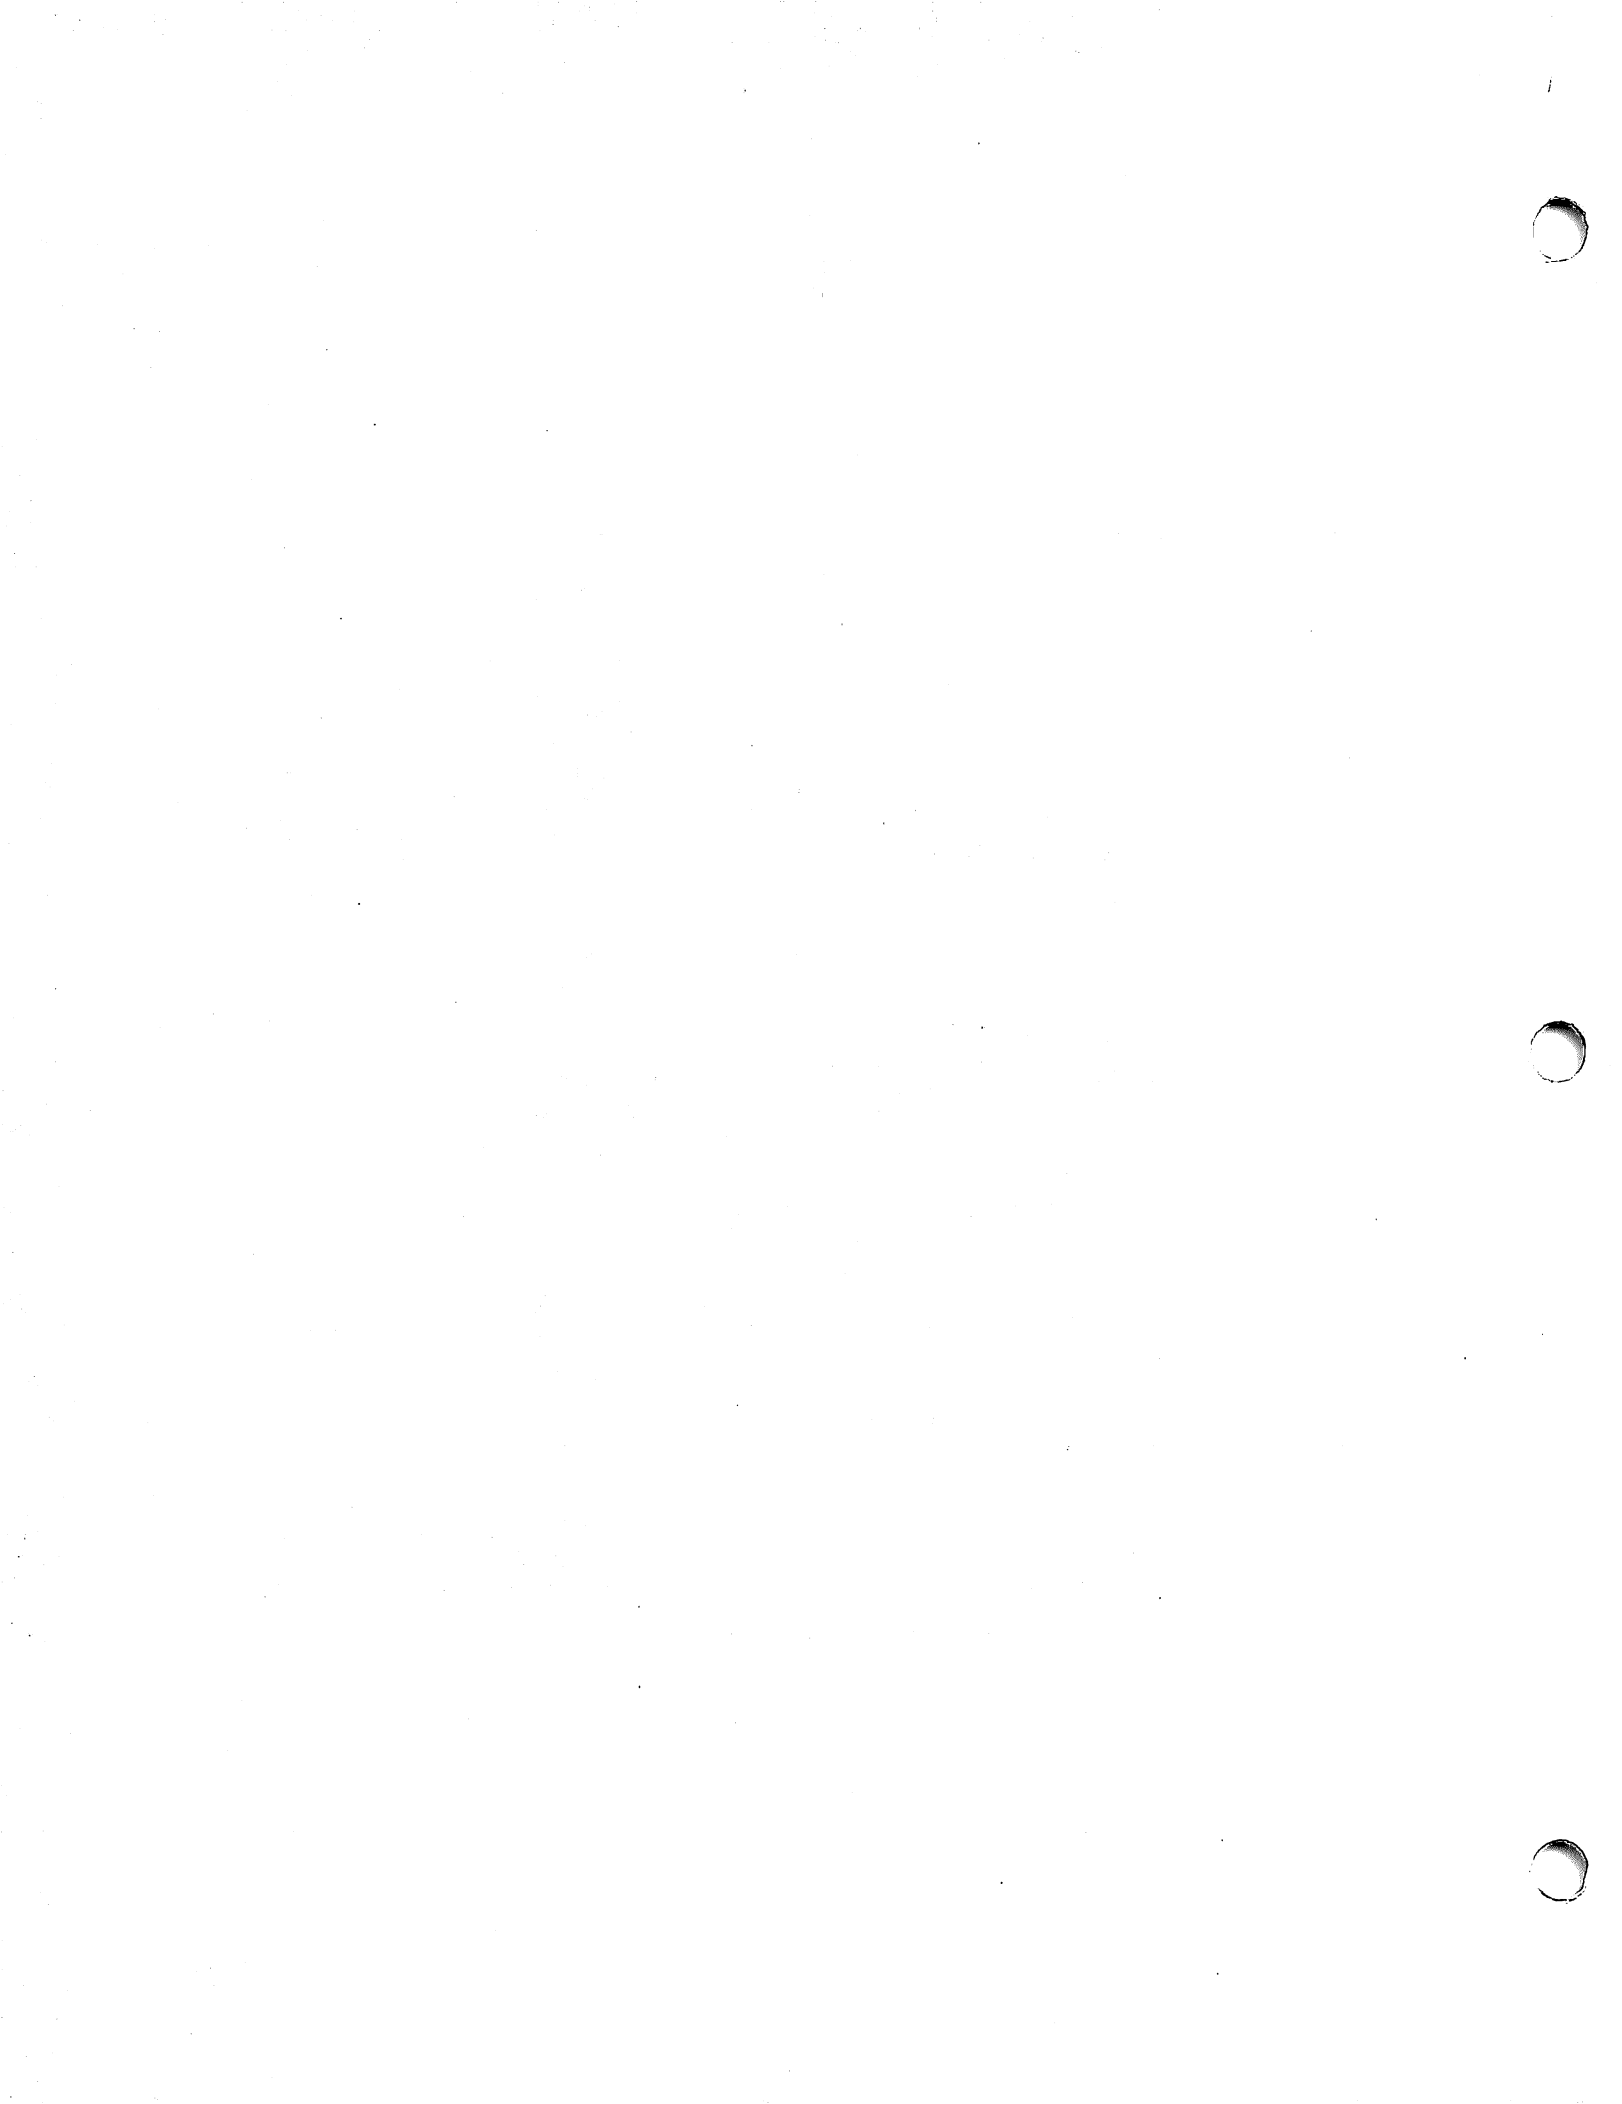
A68 Assembler Reference Manual

August 15, 1983

PLEXUS COMPUTERS, INC.
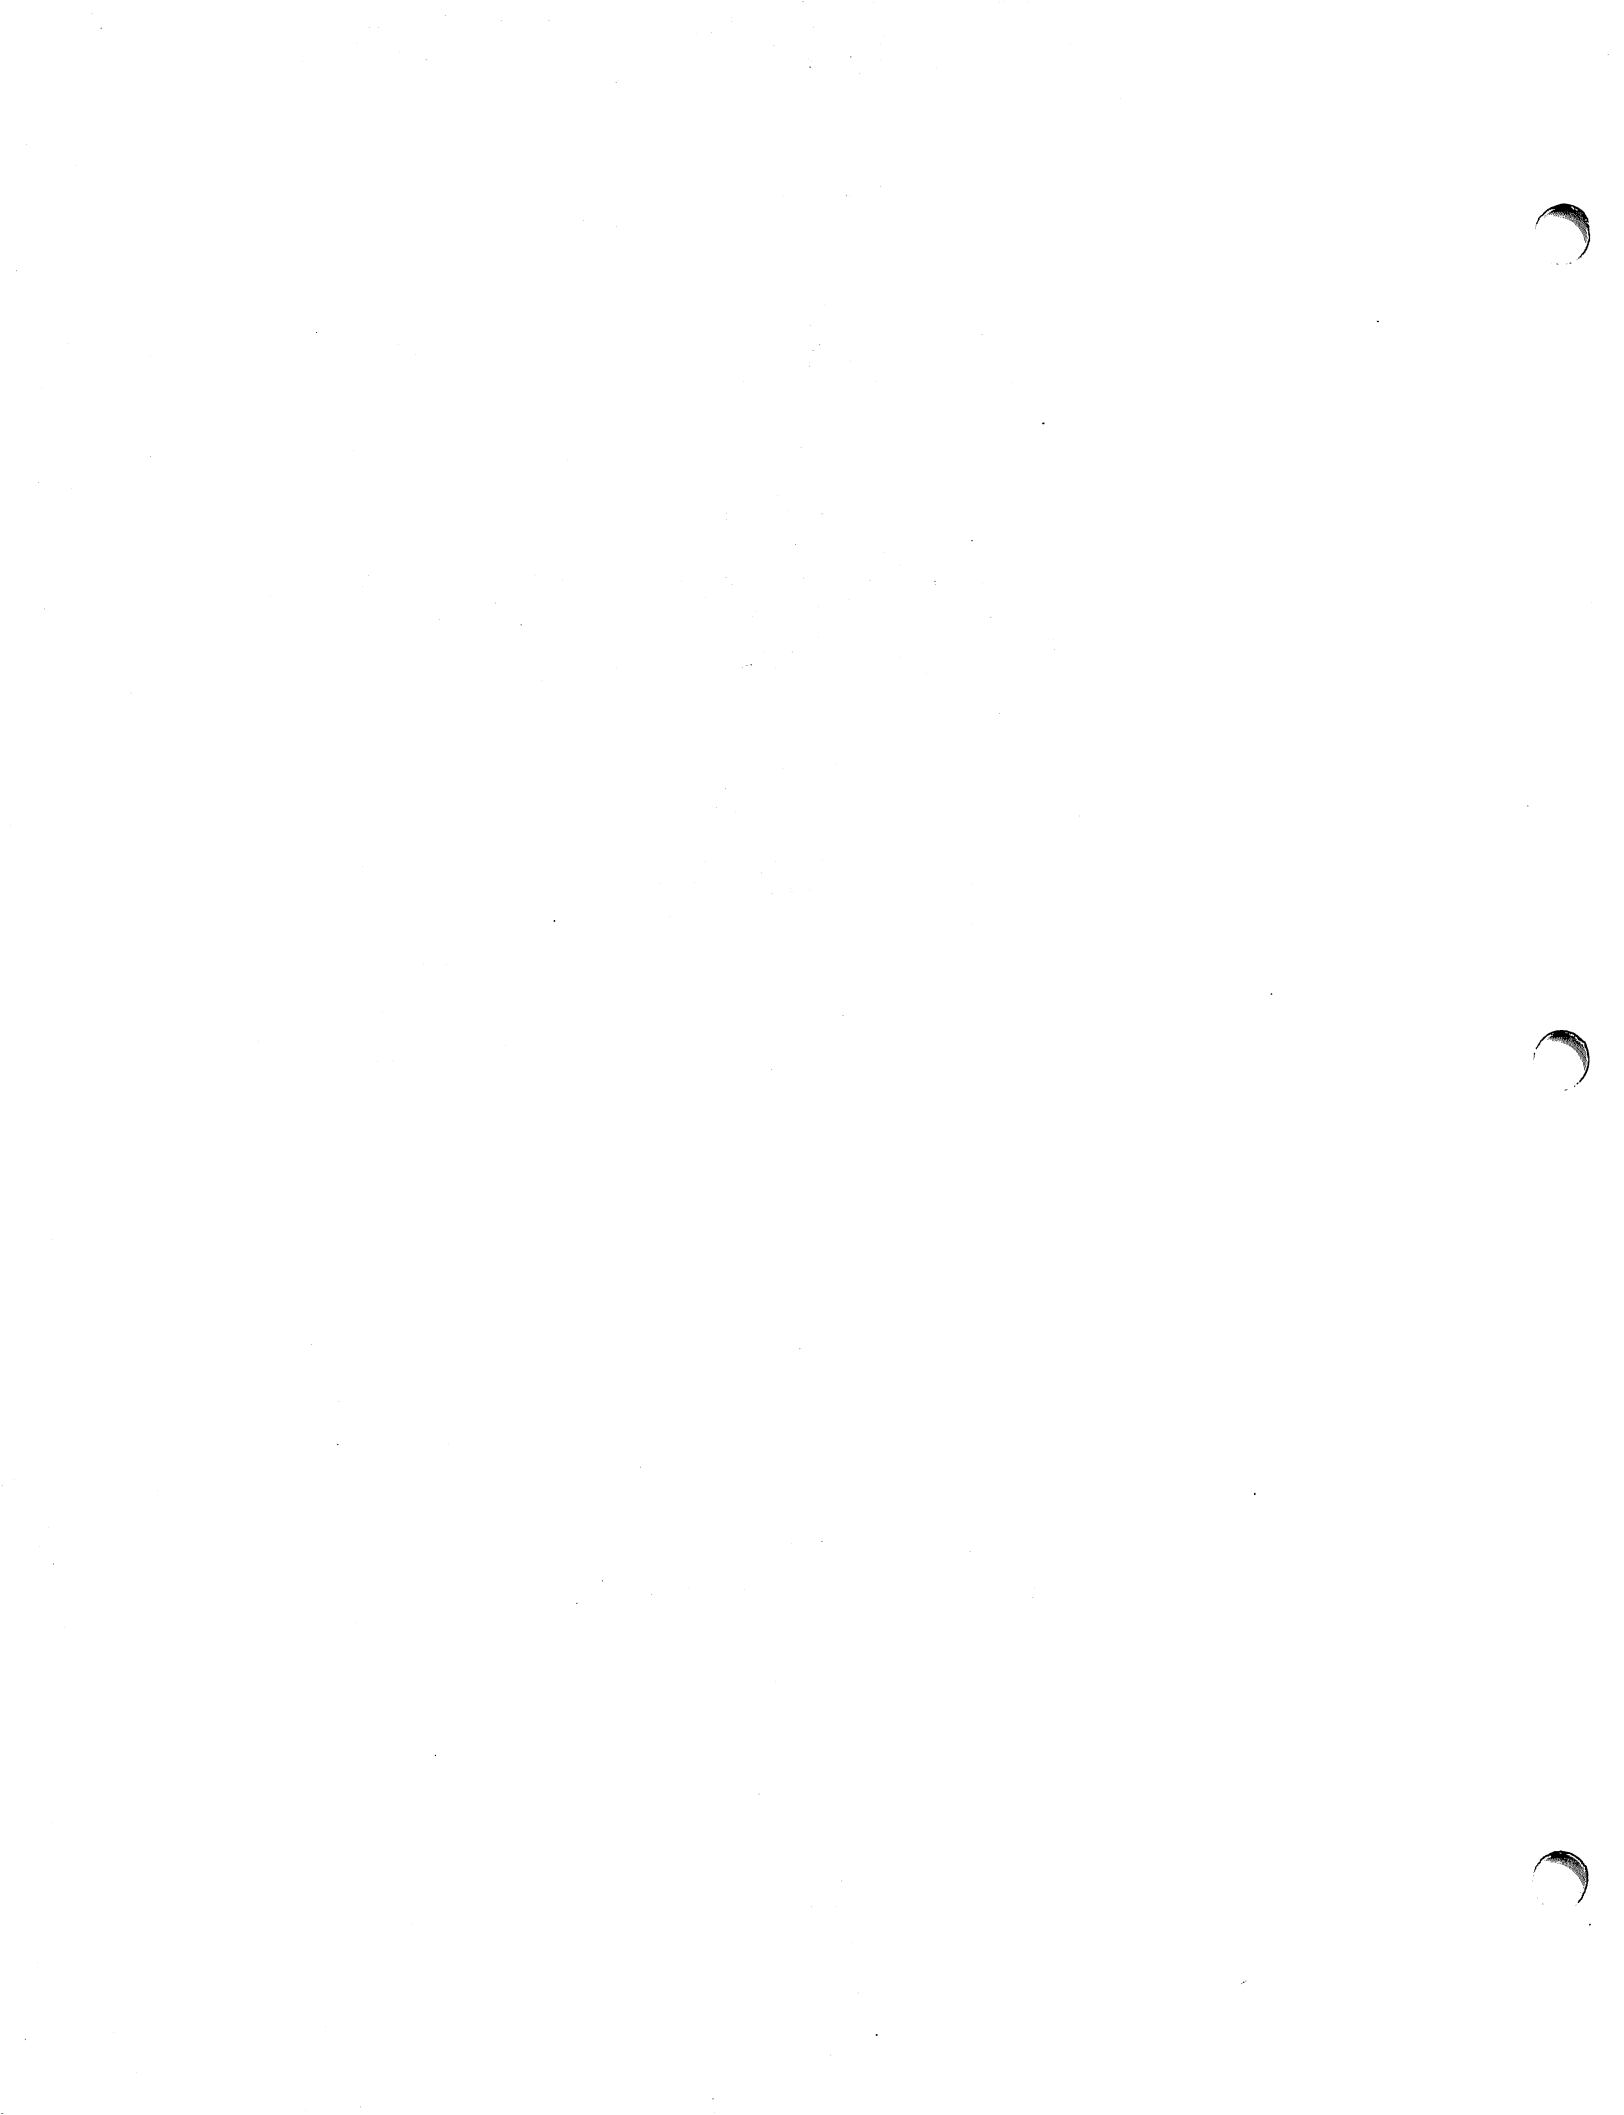
2230 Martin Ave.

Santa Clara, CA 95050
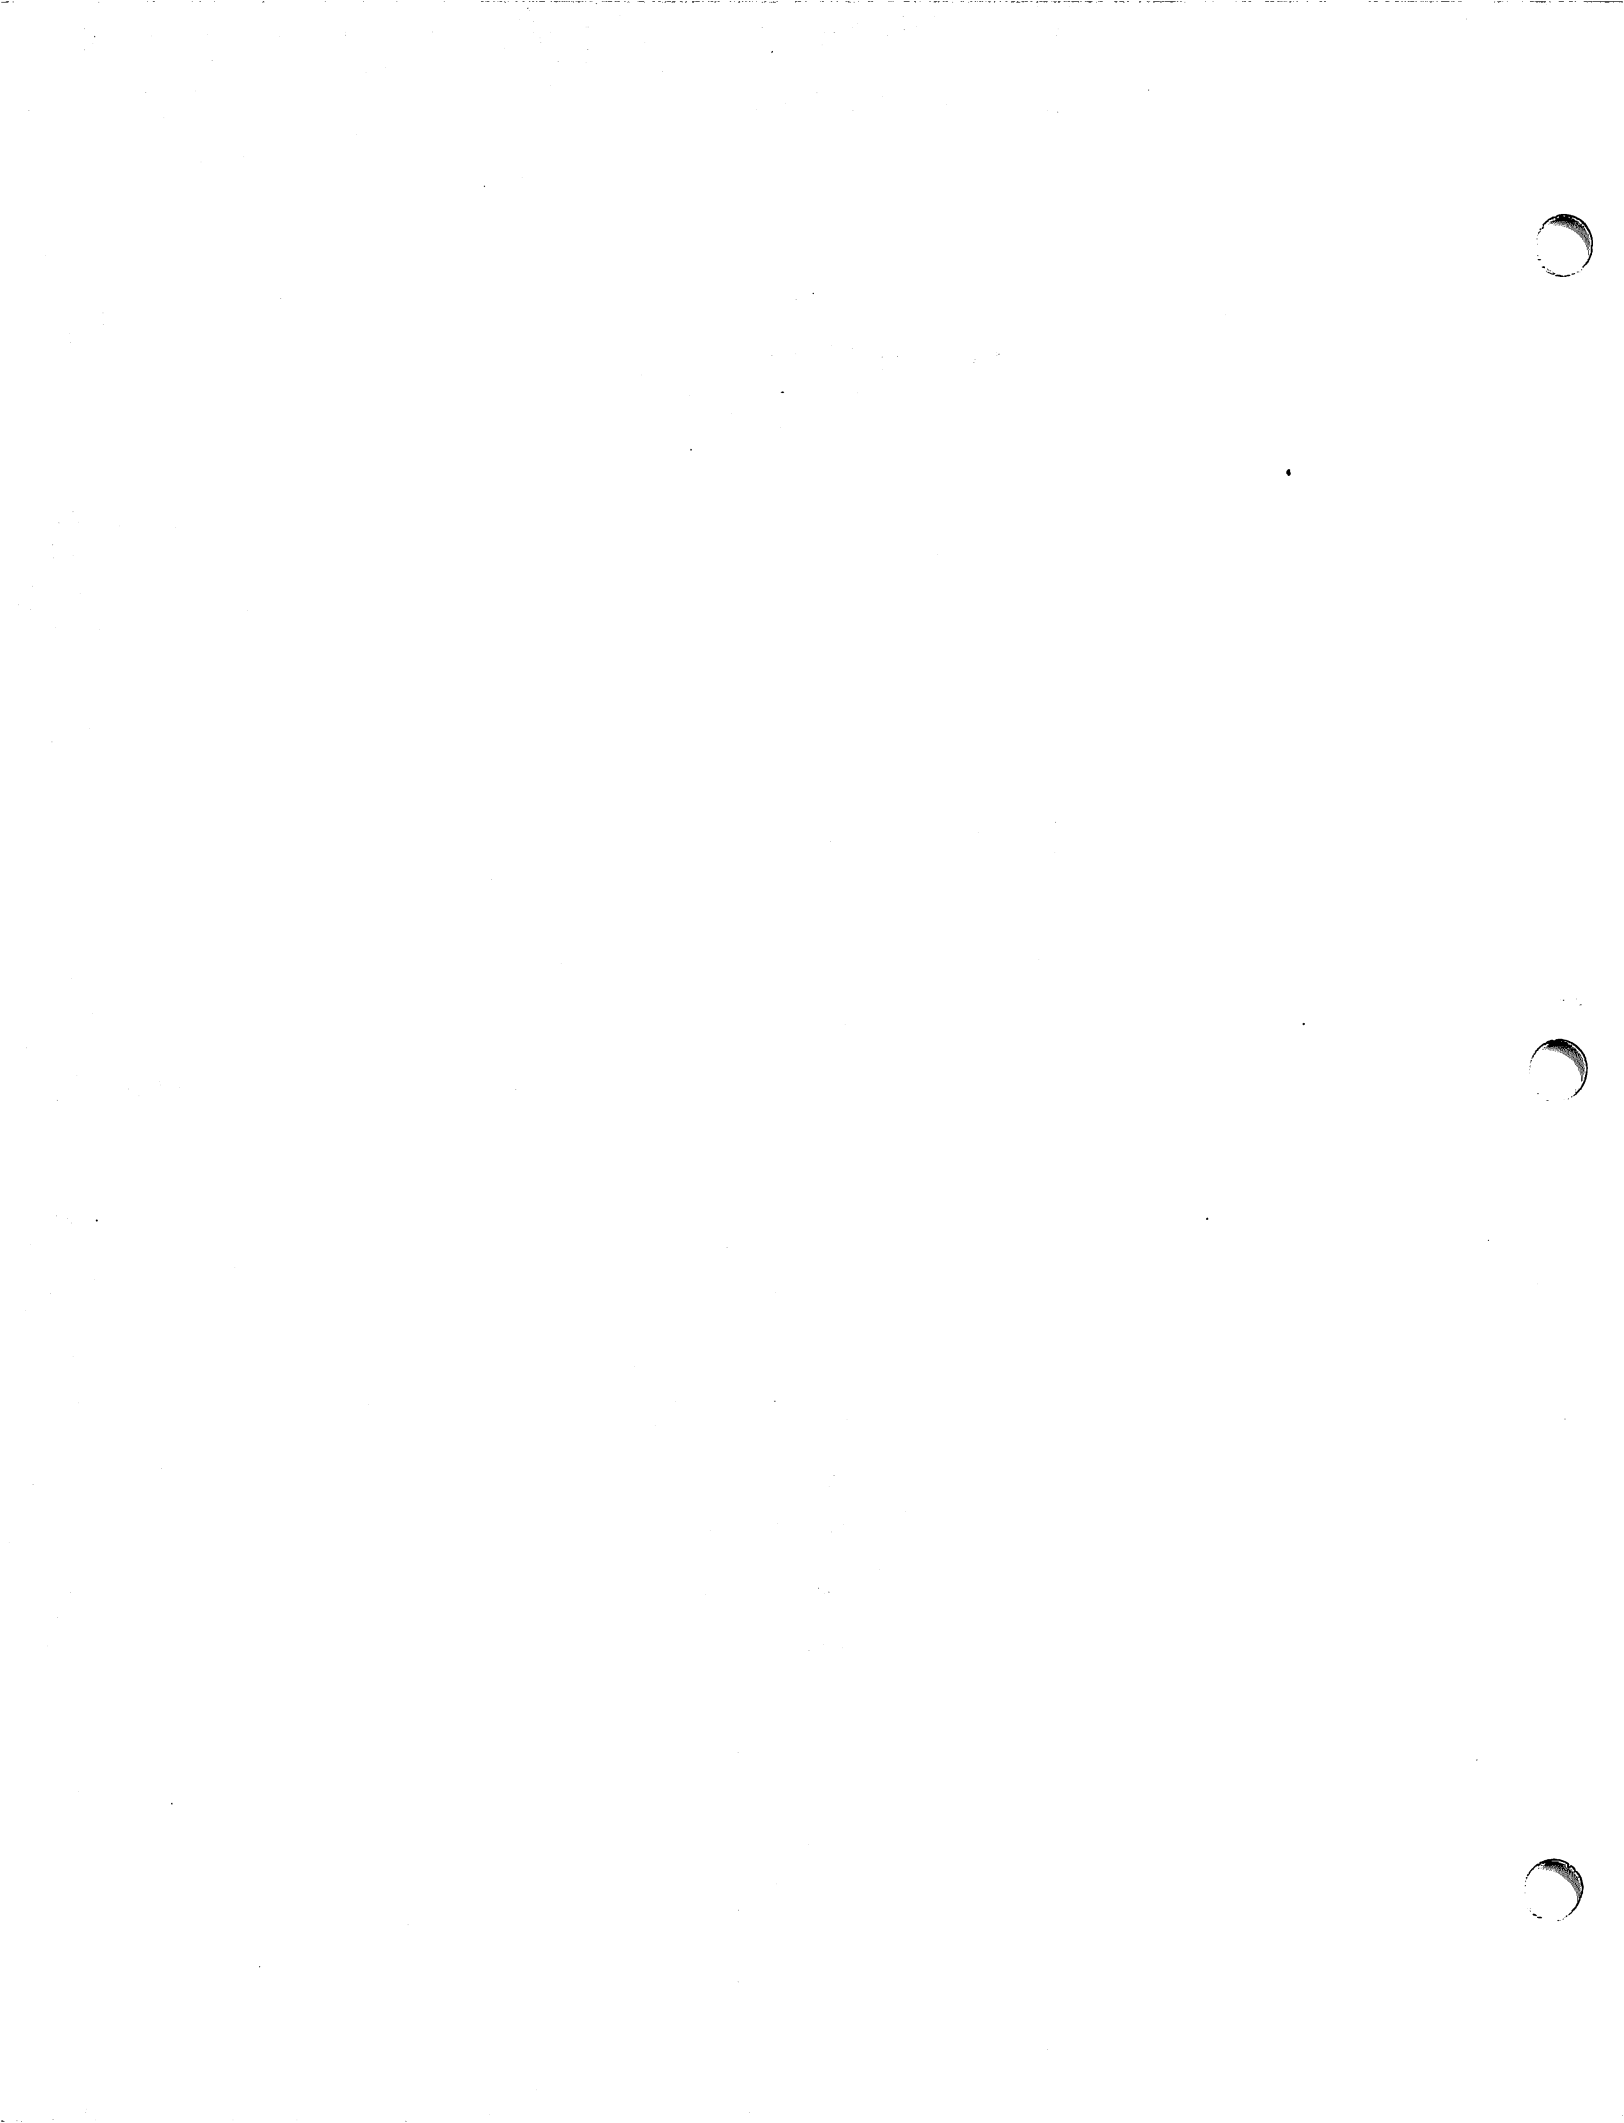
408/988-1755

Printed in the United States of America

CONTENTS

1. INTRODUCTION

This document is based on the *A68 Assembler Reference Manual* from the Massachusetts Institute of Technology Laboratory of Computer Science. It describes the syntax and usage of the **a68** assembler for the Motorola 68000 microprocessor. The basic format of **a68** is loosely based on the Digital Equipment Corp Macro-11 assembler described in DEC's publication DEC-11-0MACA-A-D. **A68** also contains elements of the UNIX* **as** assembler. The instruction mnemonics and effective address format are derived from a Motorola publication on the 68000, the *MACSS MC68000 Design Specification Intruction Set Processor* dated June 30, 1979.

Sections 2-3 of this document describe the general form of **a68** programs; section 4 defines symbols and expressions for **a68**; section 5 describes instruction mnemonics and addressing modes; section 6 describes the pseudo-ops supported by the assembler; and section 7 describes the error codes generated. For instructions on how to operate the assembler from UNIX, consult the *Plexus Sys3 UNIX Programmer's Manual -- vol 1* entry on **a68**.

2. NOTATION

The notation used in this document is a modified BNF similar to that used in the MULTICS PL/I Language Manual. The operators of the BNF in order of decreasing precedence are:

Repetition	x ...	Denotes one or more occurrences of x.
Juxtaposition	xy	Denotes an occurrence of x followed by an occurrence of y.
Alternation	x|y	Denotes an occurrence of x or y but not both.

Brackets and braces define the order of expression interpretation. A subexpression enclosed in brackets is optional. That is,

[x]	denotes zero or one occurrence of x.
{x|y}z	denotes an x or a y, followed by a z.

1

3. SOURCE PROGRAM FORMAT

3.1 Program Sections

As in UNIX, programs to **a68** are divided into three sections: *text*, *data* and *bss*. The normal interpretation of these sections is: instruction space, initialized data space and uninitialized data space, respectively. These three sections are equivalent as far as **a68** is concerned with the exception that no instructions or data are output for the *bss* section although its size is computed and its symbol values output.

In the first pass of the assembly, **a68** maintains a separate location counter for each section. Thus for code like:

```
.text           | pseudo-op that changes the program section
sum: movw d1,d2
.data           | pseudo-op that changes the program section
hello: .long 27
.text
jnk: addw d2,d1
.data
myst: .byte 4
```

in the output, *sum* will immediately precede *jnk* and *hello* will immediately precede myst. At the end of the first pass, **a68** rearranges all the addresses so that the sections will be output in the following order: *text*, *data* and *bss*. The resulting output file is an executable image file with all addresses correctly resolved, with the exception of undefined *.globls* and *.comms*. For more information on the format of the output file, see **a.out**(5).

3.2 A68 Statement Format

An **a68** program consists of a series of statements, each of which occupies exactly one line: in other words, a sequence of characters followed by the *newline* character. Form feed, ascii ^L, also serves as a line terminator. **A68** does not allow multiple statements on a single line or continuation lines.

The format of an **a68** assembly language statement is:

[LabelField :] opcode [OperandField] [|comment]

There are three exceptions to this rule:

1. Blank lines are permitted.

2. A statement may contain only a *LabelField*. The label defined in this field has the same value as if it were defined in the *LabelField* of the next statement in the program. For example, the two statements

    ```
    sea:
    movw d1,d2
    ```

    are equivalent to the single statement

    ```
    sea: movw d1,d2
    ```

3. A line may consist of only the comment field. For example, the two statements below are allowed:

    ```
    | This is a comment field.
    | So is this
    ```

In general, blanks or tabs are allowed anywhere in a statement. For example, multiple blanks are allowed in the *OperandField* to separate symbols from operators. Blanks are meaningful only when they occur in a character string (e.g. as the operand of an *.ascii* pseudo-op). At

least one blank or tab must appear between the opcode and the *OperandField* of a statement.

3.3 Label Fields

A label is a user-defined symbol. It is assigned the value of the current location counter and entered into the assembler's symbol table. The value of the label may be either absolute or relocatable; in the latter case, the absolute value of the symbol is assigned when the program is linked via **ld**(1).

A label is a symbolic means of referring to a specific location within a program. If a label is present, it occurs first in a statement and must be terminated by a colon. The collection of label definitions in a statement is called the *LabelField*.

The format of a *LabelField* is:

> symbol: [symbol:] . . .

For example:

```
start:
sea: bar:        | Multiple symbols
7$:              | A local symbol (see 4.5 below)
```

3.4 Opcode Fields

The *OpcodeField* of an assembly language statement identifies the statement as either a machine instruction, or an assembler directive. One or more blanks (or tabs) must separate the *OpcodeField* from the *OperandField* in a statement. No blanks are necessary between *LabelField*s and *OpcodeField*s, but they are recommended to improve readability of the program.

A machine instruction is indicated by an instruction mnemonic. The assembly language statement is intended to produce a single executable machine instruction. The operation of each instruction is described in the manufacturer's user manual. Some conventions used in **a68** for instruction mnemonics are described in section 4 and a complete list of the instructions is presented in the Appendix.

An assembler directive, or pseudo-op, performs some function during the assembly process. It does not produce any executable code, but it may assign space in a program for data.

3.5 Operand Fields

Note **a68** makes a distinction between *OperandField* and *operand*. Several machine instructions and assembler directives require two or more arguments, and each of these is referred to as an *operand*. In general, an *OperandField* consists of zero or more operands, and in all cases, operands are separated by a comma. In other words, the format for an *OperandField* is:

> [operand [, operand] . . .]

The format of the *OperandField* for machine instruction statements is the same for all instructions, and is described in Table 5-1. The format of the *OperandField* for assembler directives depends on the directive itself, and is included in the directive's description in section 6 of this manual.

3.6 Comment Field

The comment character in **a68** is the vertical bar, (|), not the semicolon, (;). Use of the semicolon as a comment character will result in an "Invalid Operator" error.

The comment field consists of all characters on a source line following and including the comment character. These characters are ignored by the assembler. Any character may appear in the comment field, with the obvious exception of the *newline* character, which starts a new line.

4. SYMBOLS AND EXPRESSIONS

An **a68** command line is typically composed of an opcode and an *expression*, or just an expression. An **a68** expression may be composed of symbols, numbers, terms, and other expressions.

4.1 Expressions

Expressions are combinations of *terms* joined together by binary operators. A term may be one of the following:

1. A number.

2. A symbol.

3. A term preceded by a unary operator. For example, both "sum" and "~sum" may be considered to be terms. Multiple unary operators are allowed; e.g. " -- A" has the same value as "A".

An expression is always evaluated to a 32-bit value. If the instruction calls for only one byte, (e.g., the *.byte* pseudo-op), then the low-order byte is used.

Expressions are evaluated left to right. No binary operators have precedence over each other. Thus "1+2*3" evaluates to 9, not 7. Unary operators have precedence over binary operators since they are considered part of a term, and both terms of a binary operator must be evaluated before the binary operator can be applied.

A missing expression or term is interpeted as having a value of zero. In this case, an "Invalid expression" error will be generated. An "Invalid Operator" error means that a valid end-of-line character or binary operator was not detected after the assembler processed a term. In particular, this error will be generated if an expression contains a symbol with an illegal character, or if an incorrect comment character was used.

Any expression, when evaluated, is either absolute, relocatable, or external:

a. An expression is *absolute* if its value is fixed by the second assembler pass. An expression whose terms are constants is absolute, e.g., "F1 = 32 + 56". So is an expression whose terms are symbols whose values are constants via direct assignment statements. For example, if F1 and F2 receive constant values in direct assignment statements (see 4.2.1), the expression "F1 = F2" is absolute. A relocatable expression minus a relocatable term, where both items belong to the same program section, is also absolute.

b. An expression is *relocatable* if its value is fixed relative to a base address, but will have an offset value when it is linked, or loaded into core. All labels of a program defined in relocatable sections are relocatable terms, and any expression which contains them must only add or subtract constants to their value. For example, assume the symbol "sum" was defined in a relocatable section of the program. Then the following demonstrates the use of relocatable expressions:

sum	relocatable
sum+5	relocatable
sum*2	Not relocatable (error)
2-sum	Not relocatable (error), since the expression cannot be linked by adding *sum*'s offset to it.
sum-jnk	Absolute, since the offsets added to *sum* and *jnk* cancel each other out.

An expression can contain no more than one relocatable element.

4

c. An expression is *external* (or *global*) if it contains an external symbol not defined in the current program. The same restrictions on expressions containing relocatable symbols apply to expressions containing external symbols. Exception: the expression "sum-jnk" where both *sum* and *jnk* are external symbols is not allowed.

4.2 Symbols

A symbol consists of a sequence of characters, with the following restrictions:

1. Valid characters include A-Z, a-z, 0-9, period (.), underscore (_), and dollar sign ($).

2. The first character must not be numeric.

All characters are significant and are checked in comparisons with other symbols. Upper and lower cases are distinct, so "One" and "one" are separate symbols.

A symbol is declared when the assembler recognizes it as a symbol of the program. A symbol is defined when a value is associated with it. With the exception of symbols declared by a *.globl* directive, all symbols are defined when they are declared. A label symbol (which represents an address of the program) may not be redefined; all other symbols are allowed to receive a new value.

There are several ways to declare a symbol:

1. As the label of a statement (See section 2.1).

2. In a direct assignment statement.

3. As an external symbol via the *.globl* directive.

4. As a common symbol via the *.comm* directive.

5. As a local symbol.

Cases 2-5 are described in separate subsections below.

4.2.1 Direct Assignment Statements

A direct assignment statement assigns the value of an arbitrary expression to a specified symbol. The format of a direct assignment statement is:

symbol = expression

Examples of valid direct assignments are:

vect_size = 4
vectora = 0xFFFE
vectorb = vectora - vect_size
CRLF = 0x0D0A

Only one symbol may be assigned in a single statement.

Any symbol defined by direct assignment may be redefined later in the program, in which case its value is the result of the last such statement. A local symbol may be defined by direct assignment, though this doesn't make much sense. Label or register symbols may not be redefined.

If the expression is absolute, then the symbol is also absolute, and may be treated as a constant in subsequent expressions (see section 4.1). If the expression is relocatable, however, then the symbol is also relocatable, and it is considered to be declared in the same program section as the expression.

If the expression contains an external symbol, then the symbol defined by the = statement will also be considered external. For example:

```
.globl x     | x is declared as external symbol
sum = x              | sum becomes an external symbol
```

assigns the value of x (zero if it is undefined) to sum and makes sum an external symbol. External symbols may be defined by direct assignment.

4.2.2 Register Symbols

Register symbols are symbols that represent registers in the machine. They are defined in the source descriptor file for a machine in the pre-assembly code portion of the file. This portion consists of source statements that are assembled before every source program for the machine.

The following symbols are register symbols.

```
d0  d1  d2  d3  d4  d5  d6  d7
a0  a1  a2  a3  a4  a5  a6  a7
sp  pc  cc  sr  usp
```

where

d0-d7	Data registers.
a0-a7	Address registers.
sp	Stack pointer.
pc	Program counter.
cc	Condition code.
sr	Status register.
usp	User mode stack pointer.

4.2.3 External Symbols

A program may be assembled in separate modules, and then linked together to form a single program (see ld(1)). External symbols may be defined in each of these separate modules. A symbol which is declared (given a value) in one module may be referenced in another module by declaring the symbol to be external in both modules. There are two forms of external symbols: those defined with the .globl and those defined with the .comm directive.

External symbols are declared with the .globl assembler directive. The format is:

```
.globl symbol [, symbol] . . .
```

For example, the following statements declare the array TABLE and the routine SRCH to be external symbols:

```
.globl TABLE,SRCH
TABLE: .=.+20
SRCH: movl #TABLE,a0
etc . . .
```

External symbols are only declared to the assembler. External symbols must be defined (i.e. given a value) in some other statement by one of the methods mentioned above. They need not be defined in the current program; in this case they are flagged as "undefined" in the symbol table. If they are undefined, they are considered to have a value of zero in expressions.

The other form of external symbol is defined with the .comm directive. These statement reserve storage in the bss section similar to FORTRAN common areas. The format of the statement is:

```
.comm name, ConstantExpression
```

which causes a68 to declare the name as a common symbol with a value equal to the ConstantExpression. For the rest of the assembly this symbol will be treated as though it was an undefined global. a68 does not allocate storage for common symbols; this task is left to the

loader. The loader will compute the maximum size of for each common symbol which may appear in several load modules, allocates storage for it in the final *bss* section and resolves linkages.

4.2.4 Local Symbols

Local symbols provide a convenient means of generating labels for branch instructions, etc. Use of local symbols reduces the possibility of multiply-defined symbols in a program, and separates entry point symbols from local references, such as the top of a loop. Local symbols cannot be referenced by other object modules.

Local symbols are of the form n$ where n is any integer. The following are valid local symbols

    1$
    27$
    394$

A local symbol is defined and referenced only within a single "local symbol block" (*lsb*). A new local symbol block is entered when:

1.   a label is declared; or,

2.   a new program section is entered.

Local symbols which have the same name but which appear in different local symbol blocks do not conflict.

4.3 Constants

All constants are considered absolute quantities when appearing in an expression.

4.3.1 Numeric Constants

An **a68** numeric constant is a sequence of digits. *A68* interprets integers as octal, hex, or decimal according to the following conventions:

octal	octal numbers begin with 0.
hex	hex numbers begin with 0x or 0X.
decimal	all other numbers.

4.4 Operators

4.4.1 Unary Operators

There are two unary operators in **a68**:

Operator   Function

-               unary minus

~   logical negation

4.4.2 Binary Operators
Binary operators in **a68** include:

Operator	Description	Example
+	Addition	"3+4" evaluates to 7
-	Subtraction	"3-4" evaluates to -1 or 0xFFFFFFFF
*	Multiplication	"4*3" evaluates to 12

Each operator is assumed to work on a 32-bit number.

4.5 Special Symbol -- Assembly Location Counter

The assembly location counter is the period character, '.'; hence its name "dot". When used in the operand field of any statement, dot represents the address of the first byte of the statement. Even in assembly directives, it represents the address of the start of the directive. A dot appearing as the third argument in a .byte instruction has the value of the address where the first byte was loaded; this address is not updated "during" the pseudo-op.

For example,

Ralph: movl .,a0 |load value of this instruction into a0

At the beginning of each assembly pass, the assembler clears the location counter. Normally, consecutive memory locations are assigned to each byte of generated code. However, the location where the code is stored may be changed by a direct assignment altering the location counter:

. = expression

This expression must not contain any forward references, and must not change from one pass to another. Storage area may also be reserved by advancing dot. For example, if the current value of dot is 1000, the direct assignment statement:

Table: .=.+100

would reserve 100 (decimal) bytes of storage, with the address of the first byte as the value of *Table*. The next instruction would be stored at address 110.

5. INSTRUCTIONS AND ADDRESSING MODES

This section describes the conventions used in **a68** to specify instruction mnemonics and addressing modes.

5.1 Instruction Mnemonics

The instruction mnemonics used by **a68** are described in the previously mentioned Motorola manual with a few variations. Most of the 68000 instructions can apply to byte, word on long operands, so in **a68** the normal instruction mnemonic is suffixed with *b*, *w*, or *l* to indicate which length operand was intended. For example, there are three mnemonics for the *or* instruction: *orb*, *orw* and *orl*. Opcodes for instructions with unusual opcodes may have additional suffixes; thus in addition to the normal *add* variations, there also exist: *addqb*, *addqw* and *addql* for the *add quick* instruction.

There are two kinds of branch instructions: byte and word. In **a68**, the byte (i.e., short) version is specified by appending the suffix *s* to the basic mnemonic as in *beq* and *beqs*.

In addition to the instructions which explicitly specify the instruction length, **a68** supports extended branch instructions, whose names are generally constructed by replacing the *b* with j. If the operand of the extended branch instruction is a simple address in the current segment, and the offset to that address is sufficiently small, **a68** will automatically generate the corresponding short branch instruction. If the offset is too large for a short branch, but small enough for a branch, then the corresponding branch instruction is generated.

A "complex" address is either an address which specifies other than normal mode addressing, or a relocatable expression containing more than one relocatable symbol. Thus, if *a*, *b* and *c* are symbols in the current segment, then the expression *a+b-c* is relocatable and complex. If the operand references an external address or is complex, then the extended branch instruction is implemented either by a *jmp* or *jsr* (for *ja* or *jbsr*), or by a conditional branch (with the sense of the conditional inverted) around a *jmp* for the extended conditional branches.

Note that **a68** is not optimal for extended branch instructions whose operand addresses the next instruction. The optimal code is no instruction at all, but **a68** currently retains insufficient information to make this optimization. The difficulty is that if **a68** decides to just eliminate the instruction, the address of the next instruction will be the same as the address of the (nonexistent) extended branch instruction. This instruction will then look like a branch to the current location, which would require an instruction to be generated. The object code improver *c2* (invoked by the "-O" flag to **cc**) removes such branches.

The algorithm used by **a68** for deciding how to implement extended branch instructions is described in "Assembling Code for Machines with Span-Dependent Instruction," by Thomas G. Szymanski in *Communications of the ACM*, Volume 21, Number 4, April 1978.

The Appendix gives a complete list of the instruction opcodes.

5.2 Addressing Modes

The following table describes the addressing modes recognized by **a68**. In this table *an* refers to an address register, *dn* refers to a data register, *Ri* to either a data or an address register, *d* to a displacement, which, in **a68** is a constant expression, and *xxx* to a constant expression. Certain instructions, particularly *move*, accept a variety of special registers including *sp*, the stack pointer (which is equivalent to *a7*), the status register (*sr*), the condition codes (*cc*) of the status register, the user mode stack pointer (*usp*), and the program counter (*pc*).

Mode	Notation	Example
Register	an,dn,sp,pc,cc,sr,usp	movw a3,d2
Register Indirect	an@	movw a3@,d2
Postincrement	an@+	movw a3@+,d2
Predecrement	an@-	movw a3@-,d2
Displacement	an@(d)	movw a3@(24),d2
Word Index	an@(d, Ri:W)	movw a3@(16, d2:W),d3
Long Index	an@(d, Ri:L)	movw a3@(16, d2:L),d3
Absolute Short	xxx.W	movw 14.W,d2
Absolute Long	xxx.L	movw 14.L,d2
PC Displacement	pc@(d)	movw pc@(20),d3
PC Word Index	pc@(d, Ri:W)	movw pc@(14, d2:W),d3
PC Long Index	pc@(d, Ri:L)	movw pc@(14, d2:L),d3
Normal	sun	movw sun,d3
Immediate	#xxx	movw #27+3,d3

Table 5-1: Addressing Mode Formats for Machine Instruction Statements

Normal mode actually assembles as absolute long, although the value of the constant will be flagged as relocatable to the loader. The notation for these modes is derived from the Motorola notation with the exception of the colon in index mode rather than period.

The Motorola manual presents different opcodes for instructions that use the effective address as data rather than the contents of the effective address such as *adda* for "add address". **a68** does not make this distinction because it can determine the type of the operand from its form. Thus an instruction of the form:

        sun: .word 0

        ...

        addl #sun,a0

assembles to the *add address* instruction because *sun* is known to be an address.

The 68000 tends to be very restrictive in that most instructions accept only a limited subset of the address modes above. For example, the *add address* instruction does not accept a data register as a destination. **a68** tries to check all these restrictions and generates the "illegal operand" error code for instructions that do not satisfy the address mode restrictions.

6. ASSEMBLER DIRECTIVES

The following pseudo-ops are available in **a68**:

.ascii	stores character strings
.asciz	stores character strings
.byte	stores 8-bit bytes
.word	stores 16-bit words
.long	stores 32-bit longwords
.text	Text csect
.data	Data csect
.bss	Bss csect
.globl	declares external symbols
.comm	declares common symbols

6.1 .ascii .asciz

The *.ascii* directive translates character strings into their 7-bit ASCII (represented as 8-bit bytes) equivalents for use in the source program. The format of the *.ascii* directive is as follows:

        .ascii "character string"

Obviously, a newline must not appear within the character string.

The *.asciz* directive is equivalent to the *.ascii* directive with a zero byte automatically inserted as the final character of the string. Thus, when a list or text string is to be printed, a search for the null character can terminate the string.

6.2 .byte .word .long

The *.byte*, *.word* and *.long* directives reserve bytes and words, and initialize them with certain values.

The format is:

        [label:] .byte [expression] [,expression] . .

        .

        [label:] .word [expression] [,expression] . .

        .

        [label:] .long [expression] [,expression] . .

        .

For example, the first statement reserves one byte for each expression in the operand field, and initializes the value of the byte to be the low-order byte of the corresponding expression. Note that multiple expressions must be separated by commas. A blank expression is interpreted as zero, and no error is generated.

The syntax and semantics for *.word* are identical, except that 16-bit words are reserved and initialized, of course, and *.long* uses 32-bit quantities.

6.3 .text .data .bss

These statements change the "program section" where assembled code is be loaded. See section 4.8 for an example of their use.

6.4 .globl .comm

See section 4.5.

6.5 .even

This directive advances the location counter if its current value is odd. This is useful for forcing storage allocation like *.word* directives to be on word boundries.

7. ERROR CODES

If an error is detected during assembly, a message of the form:

> line_no .error_code

is output to the standard error stream.

The following *.error* codes, and their probable cause, appear below:

2   Invalid Character.
An invalid character for a character constant or character string was encountered.

3   Multiply defined symbol.
A symbol appears twice as a label, or an attempt was made to redefine a label using an = statement.

4   Symbol storage exceeded.
No more room is left in the symbol table. Assemble portions of the program separately, then bind them together.

6   Symbol length exceeded.
A symbol of more than 31 characters was encountered.

7   Undefined symbol.
A symbol not declared by one of the methods mentioned in 4.2 was encountered. This happens when an invalid instruction mnemonic is used. This also occurs when an invalid or non-printing character occurs in the statement.

8   Invalid Constant.
An invalid digit was encountered in a number.

9   Invalid Term.
The expression evaluator could not find a valid term: symbol, constant or [*expression*]. An invalid prefix to a number or a bad symbol name in an operand will generate this.

10   Invalid Operator.
Check the operand field for a bad operator.

11   Non-relocatable expression.
If an expression contains a relocatable symbol (e.g. label) then the only operations that can be applied to it are the addition of absolute expressions or the subtraction of another relocatable symbol (which produces an absolute result).

12   Invalid operand type.
      The type field of an operand is either not defined for the machine, or represents an addressing mode incompatible with the current instruction.

13   Invalid operand.
      This is a catch-all *.error*. It appears notably when an attempt is made to assign an undefined value to dot during pass 1.

14   Invalid symbol.
      If the first token on the source line is not a valid symbol (or the beginning of a comment), this is generated. Might happen if you try an implied *.word*.

15   Invalid assignment.
      An attempt was made to redefine a label with an = statement.

16   Too many labels.
      More than 10 labels and/or symbol= 's appeared on a single statement.

17   Invalid op-code.
      An opcode mnemonic was not recognized by the assembler.

18   Invalid entry point.
      The entry point of the program (declared on the *.end* statement) must be a defined label.

19   Invalid string.
      An invalid string for *.print* or *.print* was encountered. Make sure string is enclosed in double quotes.

20   Bad filename or too many levels.
      An invalid filename was given to *.insrt*, or there were more than 10 levels of nested *.insrt*s.

22   .error statement.
      A *.error* statement was encountered during pass 2.

25   Wrong number of operands.
      This is usually a warning. Check the manufacturer's assembly manual for the correct number of operands for the current instruction.

26   Line too long.
      A statement with more that 132 characters before the newline was encountered.

27   Invalid register expression.
      Any expression inside parentheses should be absolute and have the value of a register code (register symbols). This may occur if you use parentheses for anything other than the register portion of an operand.

APPENDIX - A68 OPCODES

Double Operand Instructions

addb	add
addw	"
addl	"
andb	and
andw	"
andl	"
cmpb	compare
cmpw	"
cmpl	"
eorb	exclusive or
eorw	"
eorl	"
movb	move
movw	"
movl	"
orb	inclusive or
orw	"
orl	"
subb	subtract
subw	"
subl	"

Single Operand Instructions

clrb	clear an operand
clrw	"
clrl	"
nbcd	negate decimal with extend
negb	negate binary
negw	"
negl	"
negxb	negate binary with extend
negxw	"
negxl	"
notb	logical compliment
notw	"
notl	"
st	set all ones
sf	set all zeros
shi	set high
sls	set lower or same
scc	set carry clear
scs	set carry set
sne	set not equal
seq	set equal
svc	set no overflow
svs	set on overflow
spl	set plus
smi	set minus
slt	set less than
sge	set greater or equal

sgt	set greater than
sle	set less than or equal
tas	test operand then set
tstb	test operand
tstw	"
tstl	"

Branch Instructions

bcc	branch carry clear
bccs	"
bcs	branch on carry
bscs	"
beq	branch on equal
beqs	"
bge	branch greater or equal
bges	"
bgt	branch greater than
bgts	"
bhi	branch higher
bhis	"
ble	branch less than or equal
bles	"
bls	branch lower or same
blss	"
blt	branch less than
blts	"
bmi	branch minus
bmis	"
bne	branch not equal
bnes	"
bpl	branch positive
bpls	"
bra	branch
bras	"
bsr	subroutine branch
bsrs	"
bvc	branch no overflow
bvcs	"
bvs	branch on overflow
bvss	"

Extended Branch Instructions

jcc	jump/branch carry clear
j	jump/branch on carry
jeq	jump/branch on equal
jge	jump/branch greater or equal
jgt	jump/branch greater than
jhi	jump/branch higher
jle	jump/branch less than or equal
jls	jump/branch lower or same
jlt	jump/branch less than
jmi	jump/branch minus
jne	jump/branch not equal

jpl   jump/branch positive
jra   jump/branch
jbsr   jump/branch to subroutine
jvc   jump/branch no overflow
jvs   jump/branch on overflow

Shift Instructions

aslb   arithmetic shift left
aslw   "
asll   "
asrb   arithmetic shift right
asrw
asrl
lslb   logical shift left
lslw
lsll
lsrb   logical shift right
lsrw
lsrl
rolb   rotate left
rolw
roll
rorb   rotate right
rorw
rorl
roxlb   rotate left with extend
roxlw
roxll
roxrb   rotate right with extend
roxrw
roxrl

Miscellaneous Classes

abcd   add decimal with extend
addqb   add quick
addqw
addql
addxb   add extended
addxw
addxl
bchg   test a bit and change
bclr   test a bit and clear
bset   test a bit and set
btst   test a bit
cmpmb   compare memory
cmpmw
cmpml
chk   check register against bounds
dcnt   decrement count and branch nz
divs   signed divide
divu   unsigned divide
exg   exchange registers
extw   sign extend
extl

15

jmp	jump
jsr	jump to subroutine
lea	load effective address
link	link
moveml	move multiple registers
movemw	
movepl	move peripheral
movepw	
moveq	move quick
muls	signed multiply
mulu	unsigned multiply
nop	no operation
pea	push effective address
reset	reset machine
rte	return from exception
rtr	return and restore codes
rts	return from subroutine
sbcd	subtract decimal with extend
stop	halt machine
subqb	subtract quick
subqw	
subql	
subxb	subtract extended
subxw	
subxl	
swap	swap register halves
trap	trap
trapv	trap on overflow
unlk	unlink

UNIX Z8000 ASSEMBLER MANUAL

By Craig C. Forney

PLEXUS COMPUTERS, INC.

3833 North First St.

San Jose, CA 95134

408/943-9433

CONTENTS

1. INTRODUCTION

This document describes the usage and input of **as**, the UNIX[1] Z8000 assembler. The details of the Z8000 are not described; consult Zilog's *Z8000 PLZ/ASM Assembler Language Programming Manual*. Many constructs the assembler support are identical to constructs in the C programming language[2].

2. USAGE

AS is used as follows:

   as [-l] [-m] [-n] [-o objfile] [-u] [-v] file

where the options have the following meaning:

-l   causes a listing to be created on the standard output;

-m   places a tab at the left margin of each line when the listing option -l has been selected;

-n   causes a listing to be created on the standard output, but causes a narrower listing than the -l option. This is sometimes useful for 80 column paper, but does not contain line numbers;

-u   causes all undefined symbols in the current assembly to be made undefined-external. See the **.globl** directive below;

-v   puts the assembler in verbose mode which causes the files to be listed during the first pass, along with the errors encountered.

The output of the assembler is by default placed on the file **a.out** in the current directory; the **-o** flag causes the output to be placed on the named file. If there were no unresolved external references, and no errors detected, the output file is marked executable; otherwise, if it is produced at all, it is made non-executable.

3. LEXICAL CONVENTIONS

Assembler tokens include identifiers (alternatively *symbols* or *names*), constants, and operators.

3.1 Identifiers

An identifier consists of a sequence of alphanumeric characters (including period (.), underscore (_), and tilde ( ) s alphanumeric) of which the first may not be numeric. The first 128 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.

3.2 Constants

All constants are represented internally as long values (the magnitude is representable in 31 bits; that is, is less than 2,147,483,648). Constants are represented similarly to constants in the C language.

A decimal constant consists of a sequence of digits (0-9) that begins with a non-zero digit. An octal constant consists of a sequence of octal digits (0-7) that begins with a zero digit.

---

1. UNIX is a trademark of Bell Laboratories.

2. For details, consult *The C Programming Language* by Brian W. Kernighan and Dennis M. Ritchie, Prentice-Hall, Inc., Englewood Cliffs, NJ; 1978.

Hexadecimal constants contain hexadecimal digits (0-9,A-F) and are preceded with **0X** or **0x**. Binary constants contain binary digits (0-1) and are preceded with **0B** or **0b**.

A character constant consists of from zero to four characters enclosed in single quotes. They are right justified and null padded on the left. Certain dual-character escape sequences are acceptable in place of the ASCII character to represent new-line and other non-graphics as follows:

```
\n      NL      (012)
\s      SP      (040)
\t      HT      (011)
\e      EOT     (004)
\r      CR      (015)
\a      ACK     (006)
\p      PFX     (033)
\b      BS      (010)
\ddd    ddd
\\      \
\"      "
\'      '
```

The escape \ddd consists of the backslash followed by one, two or three octal digits which are taken to specify the value of the desired character. A special case of this construct 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 is ignored.

The last three are included so that the escape character and the character constant quote and the string constant quote can be represented.

A string constant consists of a sequence of up to 128 characters enclosed in double quotes. The dual character escape sequences are acceptable in place of the ASCII character as in character constants. String constants are used with the **.ascii** and **.asciiz** pseudo-operations. (See Section 9.)

3.3 Operators

There are several single- and double-character operators (see Section 7) . They have the same meaning as the corresponding C operators.

3.4 Blanks

Blank and tab characters may be interspersed freely between tokens, but not be used within tokens (except character constants). A blank or tab is required to separate adjacent identifiers or constants not otherwise separated.

3.5 Comments

There are two types of comments: single: and multi-line comments. Both types are ignored by the assembler. The character sequence (//) introduces a single-line comment which extends through the end of the line on which it appears. The character sequence (/*) introduces a multi-line comment, which extends until the character sequence (*/) is encountered.

4. SEGMENTS

Assembler code and data fall into three segments: the text segment, the data segment, and the bss segment. The text segment is the one 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, the 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 2K 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 the text or data) is determined by the high-watter 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 *The Location Counter* and *Assignment Statements* below.

5. 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, 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, Section 4).

6. STATEMENTS

A source program is composed of a sequence of **statements**. Statements are separated either by new lines or by semicolons. There are four kinds of statements: null statements, declaration statements, assignment statements, and keyword statements.

Any kind of statement may be preceded by one or more labels.

6.1 Labels

A 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. The error is also indicated in pass 2.

6.2 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 containing only a label.

6.3 Assignment Statements

An assignment statement consists of an identifier, an equals sign (=), and a constant expression. The value and type of the expression are assigned to the identifier. The expression may be relocatable, but must be able to be evaluated at its occurrence within the source code. This requires that each element of the expression must have been previously defined, with a declaration statement, label declaration, or previous assignment statement.

As mentioned, it is permissible to assign 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.

6.4 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.

7. EXPRESSIONS

An expression is a sequence of symbols representing a value. Its constituents are identifiers, constants, and brackets.

The evaluation of expressions is identical to the evaluation of constant expressions in C. The meaning of the operators and the order of precedence is also the same as in the C language.

Although all expressions in the assembler can be considered constant expressions, there are situations where a value is needed during the first pass (e.g., the modification of the location counter), and others where it need not be evaluated until the second pass (e.g. operands of an instruction). The first type of expression will be referred to as a *pass 1 expression*, the second type as a *pass 2 expression*.

7.1 Expression Operators

The operators are, in decreasing order of precedence, as follows:

```
()
!
* / %
+ -
<<  >>
<  <=  >  >=
==  !=
&
^
|
&&
|
```

These operators behave exactly like their counterparts in the C language. See *The C Programming Language* for details. Note that expressions may be grouped by parenthesis ( ).

7.2 Types

The assembler deals with a member 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 an error (except that certain statements 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 **zld** must be used to load the assembler's output with another routine that defines the undefined reference.

absolute   An absolute symbol is 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 the 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** segment, the value of (.) is data 0.

bss   The value of the bss symbol is measured with respect to the origin of the bss segment of a program. Like text 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** segment, 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 values and type are available to the link editor so that the program may be loaded with others that reference these symbols.

other types   Each keyword known to the assembler has a type which is used to select the routine which processes the associated keyword statement. These symbols have meaning only when used as keywords.

7.3 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

The combined rules are then: if one of the operands is undefined, the result is undefined; if both operands are absolute, the result is 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.

others   It is illegal to apply these operators to any but absolute symbols.

8. PSEUDO-OPERATIONS

The keywords listed below introduce statements that 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, lowercase tokens are literals, uppercase words are substitutable.

8.1 .byte EXPRESSION [, EXPRESSION]

The **EXPRESSIONS** in the comma separated list are pass 2 expressions, and are truncated to 8 bits and assembled in successive bytes. The expressions must be absolute.

8.2 .word EXPRESSION [, EXPRESSION]

This pseudo-operation is similar to the **.byte** pseudo operations, except that the **EXPRESSIONs** are truncated to 16 bits and assembled in successive words.

8.3 .long EXPRESSION [, EXPRESSION]

This is similar to both the **.byte** and **.word** pseudo-operation, except that the **EXPRESSIONs** are not truncated and are assembled in successive **longs**.

8.4 .ascii

This pseudo-operation must be followed by an string constant as defined in Section 3.2. The ascii values are then placed in the output file.

8.5 .asciiz

This is identical to **.ascii** except that a trailing zero is placed at the end of the string.

8.6 .page

This generates a formfeed in the listing file if the list option has been selected.

8.7 .liston

This statement resumes making a listing if the -l option has also been selected.

8.8 .listoff

This statement stops the creation of a listing file until a **.liston** statement will be assembled at a word boundary.

8.9 .if EXPRESSION

The **EXPRESSION** must be absolute and a pass 1 expression. 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.

8.10 .endif

This statement marks the end of a conditionally-assembled section of code. See **.if** above.

8.11 .globl NAME [, NAME] ...

This statement makes the **names** external. If they are otherwise defined (by assignment of 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 to 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 Section 2, it is possible to force the assembler to make all otherwise undefined symbols external.

8.12 .text

8.13 .data

8.14 .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.

8.15 .org EXPRESSION

This pseudo-operation is a way of changing the location counter (.). It requires a pass 1 expression which is evaluated and used as the current location counter within the same segment. If the type of the expression is anything but absolute, an error is generated.

8.16 .comm NAME , EXPRESSION

This pseudo-operation provides a means for defining **COMMON** (as in FORTRAN).

The type of **NAME** becomes *undefined external*, and its value becomes **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 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 explicitly defined bss-segment locations.

9. MACHINE INSTRUCTIONS

Because of the rather complicated instruction and addressing structure of the Z8000, the syntax of machine instructions is varied. Although the following sections give the syntax, the processor reference manuals should be consulted on the semantics.

9.1 Addressing Modes

With the exception of immediate data and condition codes, all assembly-language operands are expressed as addresses: register, memory, and I/O addresses. The various address modes recognized by the Z8000 assembler are as follows:

- Immediate Data
- Register
- Indirect Register
- Direct Register
- Indexed Address
- Relative Address
- Based Address
- Based Indexed Address

Special characters are used in operands to identify certain of these address modes. The characters are:

- **r** preceding a word register number;

- **rh** or **rl** preceding a byte register number;

- **rr** preceding a register pair number;

- **rq** preceding a register quadruple number;

- ***** preceding an indirect-register reference;

- **$** preceding immediate data;

- **()** used to enclose the displacement part of an indexed, based, or based index address;

- **(.)** signifying the current program counter location, usually used in relative addressing.

Not every address mode can be used by every instruction. See the Zilog Z8000 processor manuals for allowable address modes for each instruction.

Note that the syntax of operands is identical to that in Zilog's PLZ/ASM assembler, except that (*) has been substituted for (@),($) for (#), and the location counter (.) for ($); the UNIX typing conventions make (@) and (#) rather inconvenient.

9.2 Instruction Set Function Summary

Z8000 instructions can be divided functionally into nine groups:

- Operations with word, byte and long word data operands are listed with their instructions mnemonics grouped together. The suffix **b** designates a long word instruction, and no suffix designates a word instruction. For some instructions, a data size is either not applicable or depends on the segmentation mode, and thus the instruction mnemonic does not have a suffix to indicate data size.

- **src** is the source operand (**srcl** and **src2** are used to distinguish between two operands in the same instruction)

- **dst** is the destination operand

- **r** is a register operand, typically used as a counter

- **num** is a number; that is, an immediate value

- **cc** is a condition code

- **flag** is any combination of the **C, Z, S, P,** and **V** status flags

- **int** is any combination of the **VI** and **NVI** interrupt control bits

Note that the instructions that are denoted with an asterisk * are privileged; that is, they can only be executed in system mode.

9.3 Load and Exchange Instructions

clr clrb	dst	Clear
ex exb	dst,src	Exchange
ld ldb ldl	dst,src	Load
lda	dst,src	Load address
ldar	dst,src	Load address relative
ldk	dst,src	Load constant
ldm	dst,src,num	Load multiple
ldr ldrb ldrl	dst,src	Load relative
pop popl	dst,src	Pop
push pushl	dst,src	Push

9.4 Arithmetic Instructions

adc adcb	dst,src	Add with carry
add addb addl	dst,src	Add
cp cpb cpl	dst,src	Compare
dab	dst	Decimal adjust
dec decb	dst,src	Decrement
div divl	dst,src	Divide
exts extsb extsl	dst	Extend sign
inc incb	dst,src	Increment
mult multl	dst,src	Multiply
neg negb	dst	Negate
sbc sbcb	dst,src	Subtract with carry
sub subb subl	dst,src	Subtract

9.5 Logical Instructions

and andb	dst,src	And
com comb	dst	Complement
or orb	dst,src	Or
test testb testl	dst	Test
tcc tccb	cc,dst	Text condition code
xor xorb	dst,src	Exclusive or

9.6 Program Control Instructions

call	dst	Call procedure
calr	dst	Call procedure relative
djnz dbjnz	r,dst	Decrement and jump if not zero
iret		Interrupt return
jp	cc,dst	Jump
jr	cc,dst	Jump relative
ret	cc	Return from procedure
sc	src	System call

9.7 Bit Manipulation Instructions

bit bitb	dst,src	Bit test
res resb	dst,src	Reset bit
set setb	dst	src
tset tsetb	dst	Test and set

9.8 Rotate and Shift Instructions

rl rlb	dst,src	Rotate Left
rlc rlcb	dst.src	Rotate left through cary
rldb	dst,src	Rotate left digit
rr rrb	dst,src	Rotate right
rrc rccb	dst,src	Rotate right through carry
rrdb	dst,src	Rotate right digit
sda sdab sdal	dst,src	Shift dynamic arithmetic
sdl sdlb sdll	dst,src	Shift dynamic logical
sla slab slal	dst,src	Shift left arithmetic
sll sllb slll	dst,src	Shift left logical
sra srab sral	dst,src	Shift right arithmetic
srl srlb srll	dst,src	Shift right logical

9.9 Block Transfer and String Manipulation Instructions

cpd cpdb	dst,src,r,cc	Compare and decrement
cpdr cpdrb	dst,src,r,cc	Compare, decrement and repeat
cpi cpirb	dst,src,r,cc	Compare, increment and repeat
cpir cpirb	dst,src,r,cc	Compare, increment and repeat
cpsd cpsdb	dst,src,r,cc	Compare string and decrement
cpsdr cpsdrb	dst,src,r,cc	Compare string, decrement and repeat
cpsi cpsib	dst,src,r,cc	Compare string and increment
cpsir cpsirb	dst,src,r,cc	Compare string, increment and repeat
ldd lddb	dst,src,r	Load and decrement
lddr lddrb	dst,src,r	Load,decrement and repeat
ldi ldib	dst,src.r	Load and increment
ldir ldirb	dst,src,r	Load, increment and repeat
trdb	dst,src,r	Translate and decrement
trdrb	dst,src,r	Translate decrement and repeat
trib	dst,src,r	Translate and increment
trirb	dst,src,r	Translate, increment and repeat
trtdb	src1,src2,r	Translate, test and decrement
trtdbd	src1,src2,r	Translate, test, decrement and repeat
trtib	src1,src2,r	Translate, test and increment
trtirb	src1,src2,r	Translate, test, increment and repeat

9.10 Input/Output Instructions

* indicates privileged instruction

in* inb*	dst,src	Input
ind* indb*	dst,src,r	Input and decrement
indr* indrb*	dst,src,r	Input, decrement and repeat
ini* inib*	dst,src,r	Input and increment
inir* inirb*	dst,src,r	Input, increment and repeat
otdr* otdrb*	dst,src,r	Output, decrement and repeat
otir* otirb*	dst,src,r	Output, increment and repeat
out* outb*	dst,src	Output
outd* outdb*	dst,src,r	Output and decrement
outi* outib*	dst,src,r	Output and increment
sin* sinb*	dst,src	Special input
sind* sindb*	dst,src,r	Special input and decrement
sindr* sindrb*	dst,src,r	Special input, decrement and repeat
sini* sinib*	dst,src,r	Special input and increment
sinir* sinrb*	dst,src,r	Special input, increment and repeat
sotdr* sotdrb*	dst,src.r	Special output, decrement and repeat
sotir* sotirb*	dst,src,r	Special output, increment and repeat

| sout*
soutb* | dst,src | Special output |
| soutd*
soutib* | dst,src,r | Special output and decrement |

* privileged instruction

9.11 CPU Control Instructions

comflg	flag	Complement flag
di*	int	Disable interrupt
ei*	int	Enable interrupt
halt*		Halt
ldctl* ldctlb*	dst,src	Load control register
ldps*	src	Load program status
mbit*		Multi-micro bit test
mreq*	dst	Multi-micro request
mres*		Multi-micro reset
mset*		Multi-micro set
nop*		No operation
resflg*	flag	Reset flag
setflg*	flag	Set flag

* privileged instruction

10. OTHER SYMBOLS

10.1 ..

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 that refers to an absolute location, the value of (..) is subtracted. Thus the value of (..) can be taken to mean the starting memory location of the program. 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 (..).

10.2 System Calls

System call names are not defined. They may be found in the file /usr/include/sys.s.

11. DIAGNOSTICS

There are two types of errors in the assembler: (1) file I/O errors and (2) source errors. File I/O errors cause the assembler to abort along with a message indicating the problem file. Source errors cause an error message to be generated along with the line number on which it occurred.

Awk — A Pattern Scanning and Processing Language
(Second Edition)

Alfred V. Aho

Brian W. Kernighan

Peter J. Weinberger

Bell Laboratories
Murray Hill, New Jersey 07974

ABSTRACT

Awk is a programming language whose basic operation is to search a set of files for patterns, and to perform specified actions upon lines or fields of lines which contain instances of those patterns. *Awk* makes certain data selection and transformation operations easy to express; for example, the *awk* program

prints all input lines whose length exceeds 72 characters; the program

NF % 2 = = 0

prints all lines with an even number of fields; and the program

{ $1 = log($1); print }

replaces the first field of each line by its logarithm.

Awk patterns may include arbitrary boolean combinations of regular expressions and of relational operators on strings, numbers, fields, variables, and array elements. Actions may include the same pattern-matching constructions as in patterns, as well as arithmetic and string expressions and assignments, if-else, while, for statements, and multiple output streams.

This report contains a user's guide, a discussion of the design and implementation of *awk*, and some timing statistics.

September 1, 1978

Awk — A Pattern Scanning and Proce sing Language
(Second Edition)

Alfred V. Aho

Brian W. Kernighan

Peter J. Weinberger

Bell Laboratories
Murray Hill, New Jersey 07974

1. Introduction

Awk is a programming language designed to make many common information retrieval and text manipulation tasks easy to state and to perform.

The basic operation of *awk* is to scan a set of input lines in order, searching for lines which match any of a set of patterns which the user has specified. For each pattern, an action can be specified; this action will be performed on each line that matches the pattern.

Readers familiar with the UNIX† program *grep*[1] will recognize the approach, although in *awk* the patterns may be more general than in *grep*, and the actions allowed are more involved than merely printing the matching line. For example, the *awk* program

    {print $3, $2}

prints the third and second columns of a table in that order. The program

    $2 ~ /A|B|C/

prints all input lines with an A, B, or C in the second field. The program

    $1 != prev { print; prev = $1 }

prints all lines in which the first field is different from the previous first field.

1.1. Usage

The command

    awk program [files]

executes the *awk* commands in the string program on the set of named files, or on the standard input if there are no files. The statements can also be placed in a file pfile, and executed by the command

---

    awk −f pfile [files]

1.2. Program Structure

An *awk* program is a sequence of statements of the form:

    pattern   { action }
    pattern   { action }
    ...

Each line of input is matched against each of the patterns in turn. For each pattern that matches, the associated action is executed. When all the patterns have been tested, the next line is fetched and the matching starts over.

Either the pattern or the action may be left out, but not both. If there is no action for a pattern, the matching line is simply copied to the output. (Thus a line which matches several patterns can be printed several times.) If there is no pattern for an action, then the action is performed for every input line. A line which matches no pattern is ignored.

Since patterns and actions are both optional, actions must be enclosed in braces to distinguish them from patterns.

1.3. Records and Fields

Awk input is divided into "records" terminated by a record separator. The default record separator is a newline, so by default *awk* processes its input a line at a time. The number of the current record is available in a variable named NR.

Each input record is considered to be divided into "fields." Fields are normally separated by white space — blanks or tabs — but the input field separator may be changed, as described below. Fields are referred to as $1, $2, and so forth, where $1 is the first field, and $0 is the whole input record itself. Fields may

be assigned to. The number of fields in the current record is available in a variable named **NF**.

The variables **FS** and **RS** refer to the input field and record separators; they may be changed at any time to any single character. The optional command-line argument −F*c* may also be used to set **FS** to the character *c*.

If the record separator is empty, an empty input line is taken as the record separator, and blanks, tabs and newlines are treated as field separators.

The variable **FILENAME** contains the name of the current input file.

1.4. Printing

An action may have no pattern, in which case the action is executed for all lines. The simplest action is to print some or all of a record; this is accomplished by the *awk* command **print**. The *awk* program

    { print }

prints each record, thus copying the input to the output intact. More useful is to print a field or fields from each record. For instance,

    print $2, $1

prints the first two fields in reverse order. Items separated by a comma in the print statement will be separated by the current output field separator when output. Items not separated by commas will be concatenated, so

    print $1 $2

runs the first and second fields together.

The predefined variables NF and NR can be used; for example

    { print NR, NF, $0 }

prints each record preceded by the record number and the number of fields.

Output may be diverted to multiple files; the program

    { print $1 >"foo1"; print $2 >"foo2" }

writes the first field, $1, on the file foo1, and the second field on file foo2. The >> notation can also be used:

    print $1 >>"foo"

appends the output to the file foo. (In each case, the output files are created if necessary.) The file name can be a variable or a field as well as a constant; for example,

    print $1 >$2

uses the contents of field 2 as a file name.

Naturally there is a limit on the number of output files; currently it is 10.

Similarly, output can be piped into another process (on UNIX only); for instance,

    print | "mail bwk"

mails the output to bwk.

The variables **OFS** and **ORS** may be used to change the current output field separator and output record separator. The output record separator is appended to the output of the print statement.

Awk also provides the printf statement for output formatting:

    printf format expr, expr, ...

formats the expressions in the list according to the specification in format and prints them. For example,

    printf "%8.2f %10ld\n", $1, $2

prints $1 as a floating point number 8 digits wide, with two after the decimal point, and $2 as a 10-digit long decimal number, followed by a newline. No output separators are produced automatically; you must add them yourself, as in this example. The version of printf is identical to that used with C.[2]

2. Patterns

A pattern in front of an action acts as a selector that determines whether the action is to be executed. A variety of expressions may be used as patterns: regular expressions, arithmetic relational expressions, string-valued expressions, and arbitrary boolean combinations of these.

2.1. BEGIN and END

The special pattern BEGIN matches the beginning of the input, before the first record is read. The pattern END matches the end of the input, after the last record has been processed. BEGIN and END thus provide a way to gain control before and after processing, for initialization and wrapup.

As an example, the field separator can be set to a colon by

    BEGIN   { FS = ":" }
    ... *rest of program* ...

Or the input lines may be counted by

    END   { print NR }

If BEGIN is present, it must be the first pattern; END must be the last if used.

2.2. Regular Expressions

The simplest regular expression is a literal string of characters enclosed in slashes, like

/smith/

This is actually a complete *awk* program which will print all lines which contain any occurrence of the name "smith". If a line contains "smith" as part of a larger word, it will also be printed, as in

blacksmithing

Awk regular expressions include the regular expression forms found in the UNIX text editor *ed*[1] and *grep* (without back-referencing). In addition, *awk* allows parentheses for grouping, | for alternatives, + for "one or more", and ? for "zero or one", all as in *lex*. Character classes may be abbreviated: [a−zA−Z0−9] is the set of all letters and digits. As an example, the *awk* program

/[Aa]ho|[Ww]einberger|[Kk]ernighan/

will print all lines which contain any of the names "Aho," "Weinberger" or "Kernighan," whether capitalized or not.

Regular expressions (with the extensions listed above) must be enclosed in slashes, just as in *ed* and *sed*. Within a regular expression, blanks and the regular expression metacharacters are significant. To turn of the magic meaning of one of the regular expression characters, precede it with a backslash. An example is the pattern

/\/.*\//

which matches any string of characters enclosed in slashes.

One can also specify that any field or variable matches a regular expression (or does not match it) with the operators ~ and !~. The program

$1 ~ /[jJ]ohn/

prints all lines where the first field matches "john" or "John." Notice that this will also match "Johnson", "St. Johnsbury", and so on. To restrict it to exactly [jJ]ohn, use

$1 ~ /^[jJ]ohn$/

The caret ^ refers to the beginning of a line or field; the dollar sign $ refers to the end.

2.3. Relational Expressions

An *awk* pattern can be a relational expression involving the usual relational operators <, <=, ==, !=, >=, and >. An example is

$2 > $1 + 100

which selects lines where the second field is at least 100 greater than the first field. Similarly,

NF % 2 == 0

prints lines with an even number of fields.

In relational tests, if neither operand is numeric, a string comparison is made; otherwise it is numeric. Thus,

$1 >= "s"

selects lines that begin with an s, t, u, etc. In the absence of any other information, fields are treated as strings, so the program

$1 > $2

will perform a string comparison.

2.4. Combinations of Patterns

A pattern can be any boolean combination of patterns, using the operators || (or), && (and), and ! (not). For example,

$1 >= "s" && $1 < "t" && $1 != "smith"

selects lines where the first field begins with "s", but is not "smith". && and || guarantee that their operands will be evaluated from left to right; evaluation stops as soon as the truth or falsehood is determined.

2.5. Pattern Ranges

The "pattern" that selects an action may also consist of two patterns separated by a comma, as in

pat1, pat2   { ... }

In this case, the action is performed for each line between an occurrence of **pat1** and the next occurrence of **pat2** (inclusive). For example,

/start/, /stop/

prints all lines between **start** and **stop**, while

NR == 100, NR == 200 { ... }

does the action for lines 100 through 200 of the input.

3. Actions

An *awk* action is a sequence of action statements terminated by newlines or semicolons. These action statements can be used to do a variety of bookkeeping and string manipulating tasks.

3.1. Built-in Functions

Awk provides a "length" function to compute the length of a string of characters. This program prints each record, preceded by its length:

    {print length, $0}

length by itself is a "pseudo-variable" which yields the length of the current record; length(argument) is a function which yields the length of its argument, as in the equivalent

    {print length($0), $0}

The argument may be any expression.

Awk also provides the arithmetic functions sqrt, log, exp, and int, for square root, base *e* logarithm, exponential, and integer part of their respective arguments.

The name of one of these built-in functions, without argument or parentheses, stands for the value of the function on the whole record. The program

    length < 10 || length > 20

prints lines whose length is less than 10 or greater than 20.

The function substr(s, m, n) produces the substring of s that begins at position m (origin 1) and is at most n characters long. If n is omitted, the substring goes to the end of s. The function index(s1, s2) returns the position where the string s2 occurs in s1, or zero if it does not.

The function sprintf(f, e1, e2, ...) produces the value of the expressions e1, e2, etc., in the printf format specified by f. Thus, for example,

    x = sprintf("%8.2f %10ld", $1, $2)

sets x to the string produced by formatting the values of $1 and $2.

3.2. Variables, Expressions, and Assignments · ·

Awk variables take on numeric (floating point) or string values according to context. For example, in

    x = 1

x is clearly a number, while in

    x = "smith"

it is clearly a string. Strings are converted to numbers and vice versa whenever context demands it. For instance,

    x = "3" + "4"

assigns 7 to x. Strings which cannot be interpreted as numbers in a numerical context will generally have numeric value zero, but it is unwise to count on this behavior.

By default, variables (other than built-ins) are initialized to the null string, which has numerical value zero; this eliminates the need for most BEGIN sections. For example, the sums of the first two fields can be computed by

    { s1 += $1; s2 += $2 }
    END { print s1, s2 }

Arithmetic is done internally in floating point. The arithmetic operators are +, −, *, /, and % (mod). The C increment ++ and decrement −− operators are also available, and so are the assignment operators +=, −=, *=, /=, and %=. These operators may all be used in expressions.

3.3. Field Variables

Fields in *awk* share essentially all of the properties of variables — they may be used in arithmetic or string operations, and may be assigned to. Thus one can replace the first field with a sequence number like this:

    { $1 = NR; print }

or accumulate two fields into a third, like this:

    { $1 = $2 + $3; print $0 }

or assign a string to a field:

    { if ($3 > 1000)
        $3 = "too big"
      print
    }

which replaces the third field by "too big" when it is, and in any case prints the record.

Field references may be numerical expressions, as in

    { print $i, $(i+1), $(i+n) }

Whether a field is deemed numeric or string depends on context; in ambiguous cases like

    if ($1 == $2) ...

fields are treated as strings.

Each input line is split into fields automatically as necessary. It is also possible to split any variable or string into fields:

    n = split(s, array, sep)

splits the the string s into array[1], ..., array[n]. The number of elements found is returned. If the sep argument is provided, it is used as the field separator; otherwise FS is used as the separator.

3.4. String Concatenation

Strings may be concatenated. For example

length($1 $2 $3)

returns the length of the first three fields. Or in a print statement,

print $1 " is " $2

prints the two fields separated by " is ". Variables and numeric expressions may also appear in concatenations.

3.5. Arrays

Array elements are not declared; they spring into existence by being mentioned. Subscripts may have *any* non-null value, including non-numeric strings. As an example of a conventional numeric subscript, the statement

x[NR] = $0

assigns the current input record to the NR-th element of the array x. In fact, it is possible in principle (though perhaps slow) to process the entire input in a random order with the *awk* program

```
{ x[NR] = $0 }
END { ... program ... }
```

The first action merely records each input line in the array x.

Array elements may be named by non-numeric values, which gives *awk* a capability rather like the associative memory of Snobol tables. Suppose the input contains fields with values like apple, orange, etc. Then the program

```
/apple/   { x["apple"]++ }
/orange/  { x["orange"]++ }
END       { print x["apple"], x["orange"] }
```

increments counts for the named array elements, and prints them at the end of the input.

3.6. Flow-of-Control Statements

Awk provides the basic flow-of-control statements if-else, while, for, and statement grouping with braces, as in C. We showed the if statement in section 3.3 without describing it. The condition in parentheses is evaluated; if it is true, the statement following the if is done. The else part is optional.

The while statement is exactly like that of C. For example, to print all input fields one per line,

```
i = 1
while (i <= NF) {
    print $i
    ++i
}
```

The for statement is also exactly that of C:

```
for (i = 1; i <= NF; i++)
    print $i
```

does the same job as the while statement above.

There is an alternate form of the for statement which is suited for accessing the elements of an associative array:

```
for (i in array)
    statement
```

does *statement* with i set in turn to each element of array. The elements are accessed in an apparently random order. Chaos will ensue if i is altered, or if any new elements are accessed during the loop.

The expression in the condition part of an if, while or for can include relational operators like <, <=, >, >=, == ("is equal to"), and != ("not equal to"); regular expression matches with the match operators ~ and !~; the logical operators ||, &&, and !; and of course parentheses for grouping.

The break statement causes an immediate exit from an enclosing while or for; the continue statement causes the next iteration to begin.

The statement next causes *awk* to skip immediately to the next record and begin scanning the patterns from the top. The statement exit causes the program to behave as if the end of the input had occurred.

Comments may be placed in *awk* programs: they begin with the character # and end with the end of the line, as in

print x, y # this is a comment

4. Design

The UNIX system already provides several programs that operate by passing input through a selection mechanism. *Grep*, the first and simplest, merely prints all lines which match a single specified pattern. *Egrep* provides more general patterns, i.e., regular expressions in full generality; *fgrep* searches for a set of keywords with a particularly fast algorithm. *Sed*[1] provides most of the editing facilities of the editor *ed*, applied to a stream of input. None of these programs provides numeric capabilities, logical relations, or variables.

Lex[3] provides general regular expression recognition capabilities, and, by serving as a C program generator, is essentially open-ended in its capabilities. The use of *lex*, however, requires a knowledge of C programming, and a *lex* program must be compiled and loaded before use, which discourages its use for one-shot applications.

Awk is an attempt to fill in another part of the matrix of possibilities. It provides general regular expression capabilities and an implicit input/output loop. But it also provides convenient numeric processing, variables, more general selection, and control flow in the actions. It does not require compilation or a knowledge of C. Finally, *awk* provides a convenient way to access fields within lines; it is unique in this respect.

Awk also tries to integrate strings and numbers completely, by treating all quantities as both string and numeric, deciding which representation is appropriate as late as possible. In most cases the user can simply ignore the differences.

Most of the effort in developing *awk* went into deciding what *awk* should or should not do (for instance, it doesn't do string substitution) and what the syntax should be (no explicit operator for concatenation) rather than on writing or debugging the code. We have tried to make the syntax powerful but easy to use and well adapted to scanning files. For example, the absence of declarations and implicit initializations, while probably a bad idea for a general-purpose programming language, is desirable in a language that is meant to be used for tiny programs that may even be composed on the command line.

In practice, *awk* usage seems to fall into two broad categories. One is what might be called "report generation" — processing an input to extract counts, sums, sub-totals, etc. This also includes the writing of trivial data validation programs, such as verifying that a field contains only numeric information or that certain delimiters are properly balanced. The combination of textual and numeric processing is invaluable here.

A second area of use is as a data transformer, converting data from the form produced by one program into that expected by another. The simplest examples merely select fields, perhaps with rearrangements.

5. Implementation

The actual implementation of *awk* uses the language development tools available on the UNIX operating system. The grammar is specified with *yacc*;[4] the lexical analysis is done by *lex*; the regular expression recognizers are deterministic finite automata constructed directly from the expressions. An *awk* program is translated into a parse tree which is then directly executed by a simple interpreter.

Awk was designed for ease of use rather than processing speed; the delayed evaluation of variable types and the necessity to break input into fields makes high speed difficult to achieve in any case. Nonetheless, the program has not proven to be unworkably slow.

Table 1 below shows the execution (user + system) time on a PDP-11/70 of the UNIX programs *wc*, *grep*, *egrep*, *fgrep*, *sed*, *lex*, and *awk* on the following simple tasks:

1. count the number of lines.

2. print all lines containing "doug".

3. print all lines containing "doug", "ken" or "dmr".

4. print the third field of each line.

5. print the third and second fields of each line, in that order.

6. append all lines containing "doug", "ken", and "dmr" to files "jdoug", "jken", and "jdmr", respectively.

7. print each line prefixed by "line-number : ".

8. sum the fourth column of a table.

The program *wc* merely counts words, lines and characters in its input; we have already mentioned the others. In all cases the input was a file containing 10,000 lines as created by the command *ls* −*l*; each line has the form

−rw−rw−rw− 1 ava 123 Oct 15 17:05 xxx

The total length of this input is 452,960 characters. Times for *lex* do not include compile or load.

As might be expected, *awk* is not as fast as the specialized tools *wc*, *sed*, or the programs in the *grep* family, but is faster than the more general tool *lex*. In all cases, the tasks were about as easy to express as *awk* programs as programs in these other languages; tasks involving fields were considerably easier to express as *awk* programs. Some of the test programs are shown in *awk*, *sed* and *lex*.

References

1. K. Thompson and D. M. Ritchie, *UNIX Programmer's Manual*, Bell Laboratories (May 1975). Sixth Edition

2. B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Prentice-Hall, Englewood Cliffs, New Jersey (1978).

3. M. E. Lesk, "Lex — A Lexical Analyzer Generator," Comp. Sci. Tech. Rep. No. 39, Bell Laboratories, Murray Hill, New Jersey (October 1975).

4. S. C. Johnson, "Yacc — Yet Another Compiler-Compiler," Comp. Sci. Tech. Rep. No. 32, Bell Laboratories, Murray Hill, New Jersey (July 1975).

/

Program	Task 1	2	3	4	5	6	7	8
wc	8.6							
grep	11.7	13.1						
egrep	6.2	11.5	11.6					
fgrep	7.7	13.8	16.1					
sed	10.2	11.6	15.8	29.0	30.5	16.1		
lex	65.1	150.1	144.2	67.7	70.3	104.0	81.7	92.8
awk	15.0	25.6	29.9	33.3	38.9	46.4	71.4	31.1

Table 1. Execution Times of Programs. (Times are in sec.)

The programs for some of these jobs are shown below. The *lex* programs are generally too long to show.

AWK:

1. END {print NR}

2. /doug/

3. /ken|doug|dmr/

4. {print $3}

5. {print $3, $2}

6. /ken/   {print >"jken"}
   /doug/ {print >"jdoug"}
   /dmr/   {print >"jdmr"}

7. {print NR ": " $0}

8.   {sum = sum + $4}
   END {print sum}

SED:

1. $=

2. /doug/p

3. /doug/p
   /doug/d
   /ken/p
   /ken/d
   /dmr/p
   /dmr/d

4. /[^]* [ ]*[^]* [ ]*\([^]*\) .*/s//\1/p

5. /[^]* [ ]*\([^]*\) [ ]*\([^]*\) .*/s//\2 \1/p

6. /ken/w jken
   /doug/w jdoug
   /dmr/w jdmr

LEX:

1. %{
   int i;
   %}
   %%
   \n   i++;
   .   ;
   %%
   yywrap() {
       printf("%d\n", i);
   }

2. %%
   ^.*doug.*$ printf("%s\n", yytext);
   .   ;
   \n   ;