A/UX Text-Editing Tools

Release 3.0

Contents

2   Using TextEditor / 2-1

3   Using the `vi` Screen Editor / 3-1

Contents   **ix**

Tables

About This Guide

Welcome to *A/UX Text-Editing Tools*. This book presents detailed information on the five text editors provided by A/UX. This guide describes how to use the various text editors to create and edit text. The companion book, *A/UX Text-Processing Tools*, describes how to use UNIX® tools to format text, tables, graphs, equations, and line drawings. The two books together are the only books you'll need to design documents to suit your specific needs.

This guide contains the following chapters:

- Chapter 1, "An Overview of A/UX Text Editors"
- Chapter 2, "Using Text Editor"
- Chapter 3, "Using the `vi` Screen Editor"
- Chapter 4 ,"Using the `ex` Line Editor"
- Chapter 5, "Using the `ed` Line Editor"
- Chapter 6, "Using the `sed` Stream Editor"

Who should use this guide

This book is geared toward the person already familiar with a UNIX text editor and needing further information on its use. It helps the person unfamiliar with UNIX text editors by providing an overview of the editors. It also describes how to use the Macintosh-style editor, *TextEditor,* which is the easiest editor to learn.

What you should already know

To use TextEditor, you need to know the basics of using a Macintosh, such as double-clicking the mouse to open a file and dragging with the mouse to choose a menu command. To use the other editors, you need to know the basics of using CommandShell, the application that provides the command-line interface. CommandShell is described in *A/UX Essentials.*

How to use this guide

 If you are new to UNIX text editors, first read Chapter 1, "An Overview of A/UX Text Editors," to learn which editor best fits your needs. Then go directly to the chapter describing the editor you want to use. If you are unfamiliar with A/UX itself, first see *A/UX Essentials.*

Conventions used in this guide

A/UX guides follow specific conventions. For example, words that require special emphasis appear in specific fonts or font styles. The following sections describe the conventions used in all A/UX guides.

Keys and key combinations

Certain keys on the keyboard have special names. These modifier and character keys, often used in combination with other keys, perform various functions. In this guide, the names of these keys are in Initial Capital letters followed by SMALL CAPITAL letters.

The key names are

CAPS LOCK	DOWN ARROW (↓)	OPTION	SPACE BAR
COMMAND (⌘)	ENTER	RETURN	TAB
CONTROL	ESCAPE	RIGHT ARROW (→)	UP ARROW (↑)
DELETE	LEFT ARROW (←)	SHIFT	

Sometimes you will see two or more names joined by hyphens. The hyphens indicate that you use two or more keys together to perform a specific function. For example,

Press COMMAND-K

means "Hold down the COMMAND key and then press the K key."

Terminology

In A/UX guides, a certain term can represent a specific set of actions. For example, the word *enter* indicates that you type a series of characters on the command line and press the RETURN key. The instruction

Enter ls

means "Type ls and press the RETURN key."

Here is a list of common terms and the corresponding actions you take.

Term	Action
Click	Press and then immediately release the mouse button.
Drag	Position the mouse pointer, press and hold down the mouse button while moving the mouse, and then release the mouse button.
Choose	Activate a command in a menu. To choose a command from a pull-down menu, position the pointer on the menu title and hold down the mouse button. While holding down the mouse button, drag down through the menu until the command you want is highlighted. Then release the mouse button.
Select	Highlight a selectable object by positioning the mouse pointer on the object and clicking.
Type	Type a series of characters *without* pressing the RETURN key.
Enter	Type the series of characters indicated and press the RETURN key.

The Courier font

Throughout A/UX guides, words that appear on the screen or that you must type exactly as shown are in the Courier font.

For example, suppose you see this instruction:

Type `date` on the command line and press RETURN.

The word `date` is in the `Courier` font to indicate that you must type it.

Suppose you then read this explanation:

After you press Return, information such as this appears on the screen:

`Tues Oct 17 17:04:00 PDT 1989`

In this case, `Courier` is used to represent the text that appears on the screen.

All A/UX manual page names are also shown in the `Courier` font. For example, the entry `ls(1)` indicates that `ls` is the name of a manual page in an A/UX reference manual. See "Manual Page Reference Notation," later in this preface, for more information on the A/UX command reference manuals.

Font styles

Italics are used to indicate that a word or set of words is a placeholder for part of a command. For example,

`cat` *filename*

tells you that *filename* is a placeholder for the name of a file you want to display. For example, if you wanted to display the contents of a file named `Elvis`, you would type the word `Elvis` in place of *filename*. In other words, you would enter

`cat Elvis`

New terms appear in **boldface** where they are defined. Boldface is also used for steps in a series of instructions.

A/UX command syntax

A/UX commands follow a specific command syntax. A typical A/UX command gives the command name first, followed by options and arguments. For example, here is the syntax for the `wc` command:

`wc [-l] [-w] [-c] [`*filename*`]...`

In this example, `wc` is the command; `-l`, `-w`, and `-c` are options; and *filename* is an argument. Brackets ([ ]) enclose elements that are not necessary for the command to execute. The ellipsis (…) indicates that you can specify more than one argument. Brackets and ellipses are *not* to be typed. Also, note that each command element is separated from the next element by a space.

The following table gives more information about the elements of an A/UX command.

Element	Description
command	The command name.
option	A character or group of characters that modifies the command. Most options have the form *-option,* where *option* is a letter representing an option. Most commands have one or more options.
argument	A modification or specification of a command, usually a filename or symbols representing one or more filenames.
[ ]	Brackets used to enclose an optional item—that is, an item that is not essential for execution of the command.
…	Ellipses are used to indicate that you can enter more than one argument.

For example, the `wc` command is used to count lines, words, and characters in a file. Thus, you can enter

```
wc -w Priscilla
```

In this command line, `-w` is the option that instructs the command to count all of the words in the file, and the argument `Priscilla` is the file to be searched.

Manual page reference notation

The *A/UX Command Reference,* the *A/UX Programmer's Reference,* the *A/UX System Administrator's Reference,* the *X11 Command Reference for A/UX,* and the *X11 Programmer's Reference for A/UX* contain descriptions of commands, subroutines, and other related information. Such descriptions are known as *manual pages* (often shortened to *man pages*). Manual pages are organized within these references by section numbers.

The standard A/UX cross-reference notation is

command (section)

where *command* is the name of the command, file, or other facility; and *section* is the number of the section in which the item resides.

- Items followed by section numbers (1M) and (8) are described in the *A/UX System Administrator's Reference.*
- Items followed by section numbers (1) and (6) are described in the *A/UX Command Reference.*
- Items followed by section numbers (2), (3), (4), and (5) are described in the *A/UX Programmer's Reference.*
- Items followed by section number (1X) are described in the *X11 Command Reference for A/UX.*
- Items followed by section numbers (3X) and (3Xt) are described in the *X11 Programmer's Reference for A/UX.*

For example

```
cat(1)
```

refers to the command `cat`, which is described in Section 1 of the *A/UX Command Reference.*

You can display manual pages on the screen by using the `man` command. For example, you could enter the command

```
man cat
```

to display the manual page for the `cat` command, including its description, syntax, options, and other pertinent information. To exit a manual page, press the SPACE BAR until you see a command prompt, or type `q` at any time to return immediately to your command prompt.

For more information

To find out where you need to go for more information about how to use A/UX, see *Road Map to A/UX.* This guide contains descriptions of each A/UX guide and ordering information for all the guides in the A/UX documentation suite.
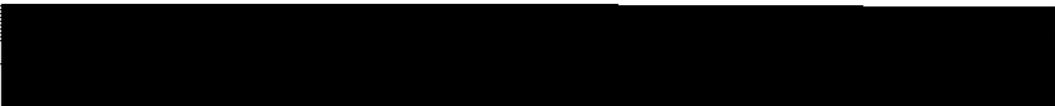
1 An Overview of A/UX Text Editors

This chapter provides a brief overview of each of the five text-editing programs available with A/UX.

What is a text editor?

A **text editor** is a program that accepts text from the keyboard, stores it in a file, and allows you to edit it. A text-editing program differs from a word-processing program in one significant way: a text-editing program operates in distinct modes, whereas a word-processing program provides a seamless integration of text-entering, -editing, and -formatting functions. With the exception of TextEditor, the editors provided with A/UX are strictly text-editing programs.

TextEditor combines features of both word-processing and text-editing programs. It operates as a word processor by allowing you to enter, edit, and format text in a seamless manner. However, it operates as a text editor by storing the formatting information in a separate file. Storing the formatting information separately allows application programs to access the file with or without the formatting information.

There are three types of text-editing programs: interactive, line, and stream. An **interactive editor** allows you to enter text and text-editing commands while you are viewing the text. A **line editor** limits you to working one line at a time on file contents. A **stream editor** is a noninteractive editor that you use to edit an existing file. Using a stream editor makes the editing process much faster, but you don't see your changes until the editor has finished with the entire document. Stream-editing commands are usually kept in a file. You call the commands from the file instead of entering them from the keyboard.

A/UX text editors

Your A/UX system has five text-editing programs: TextEditor, `vi`, `ex`, `ed`, and `sed`. These are described briefly in the sections that follow. Each of these editors is discussed in a separate chapter in this manual. Table 1-1 provides a brief overview of each editor.

TextEditor

TextEditor, the default screen editor for A/UX, is the most intuitive and easy to use of all the A/UX editors. It resembles many of the basic word-processing applications currently

available for the Macintosh computer. You open, close, save, quit, cut, paste, and perform many other basic editing tasks by choosing commands from menus. If you are new to A/UX or the UNIX® operating system, you are likely to use this editor for most of your basic editing tasks. TextEditor is the only editor that supports the use of the mouse.

Table 1-1 A/UX text editors

Editor	Main feature	Primary use
TextEditor	Mouse-based editor. This editor is the easiest to learn.	Creates new text files and edits existing ones.
vi	UNIX-standard screen editor. This editor is the most versatile.	Creates new text files and edits existing ones.
ex	Line editor. This editor includes many search, replace, and movement commands.	Teams with vi to create a powerful A/UX text-editing system.
ed	Line editor. This editor works in both Macintosh and A/UX environments.	Edits A/UX system files from the A/UX Startup application when A/UX isn't running.
sed	Stream editor. This editor allows quick, one-pass editing with filters.	Used for batch editing. Takes input from a file, not the keyboard.

The vi screen editor

The vi screen editor is the most widely used UNIX text-editing tool. The vi, or *visual,* editor was derived from another popular UNIX text-editing tool, the ex line editor program. Although the ex line editor is not a true visual editor, it has many powerful commands. Since vi was built around ex, you can access most ex commands directly from vi. Most users prefer to use vi as a front-end editor and access ex commands from ex command mode. In essence, vi works best when it is teamed with ex.

In addition to ex commands, vi has many cursor-movement commands. You can move the cursor character, word, line, sentence, paragraph, or section. You can also move the cursor to a particular character string. The vi editor does not support the use of the mouse.

The `ex` line editor

The `ex` editor provides a single-line window into the text-editing buffer and has the advantages of reduced system overhead and accessibility from a shell program. It is most commonly used in conjunction with the `vi` editor, which is an offspring of `ex` itself.

The `ex` editor is a powerful editing tool for making substitutions and giving global commands. The `ex` editor can search for a pattern and perform substitutions on any string that matches that pattern. This greatly increases the power and flexibility of substitution commands.

The `ex` editor has options that define the editing environment for the `ex` and `vi` editors (such as the margin for word wraparound on your screen, automatic indent following a line that starts with a tab character, the screen display of line numbers, a special environment for editing programs, and so on). The `ex` editor also has macro facilities that allow you to "map" complex editing sequences to a single key combination or abbreviate a long string to a short one.

All these `ex` capabilities are accessible not only from `vi` but also from shell programs using the `ex` editor alone. The `ex` line editor does not support the use of the mouse.

The `ed` line editor

The `ed` line editor provides a single-line window into the text-editing buffer. This editor devotes no time or system resources to redrawing the screen and can be an efficient way to enter text when you are working at 1200 baud or lower. In A/UX, `ed` is the only editor that allows you to work on A/UX startup files such as `inittab`, when you cancel the startup application (formerly known as `sash`). You may also use `ed` to edit a shell program since `ed` does not operate on a full screen of text.

The `sed` stream editor

The `sed` stream editor is useful for creating filters for **batch editing.** Batch editing means running a file through a series of predetermined editing commands (filters) that automatically edit the file without user supervision.

The `sed` stream editor copies the input file(s) to the standard output, performing various user-specified editing tasks (such as substituting or deleting words) on the file as it "flows" by. These tasks may be specified on the command line or, more commonly, stored in a file for repeated use.

Because of its batch nature, `sed` is extremely useful for building filters to edit or modify text without user supervision. Thus, `sed` may run in the background, allowing you to perform other tasks while the editing takes place. The `sed` stream editor is also useful for editing very large files, since it doesn't use a buffer.

The changes specified in the `sed` script (or on the command line) affect only a copy of the file, not the original file itself. The output of the `sed` command is directed to the standard output, usually your terminal screen. You may, however, redirect the output to a file or to a pipeline for further filtering by other A/UX utilities.

Using commands with A/UX text-editing programs

When you begin an editing session, everything you type is interpreted as a command. Each A/UX editor has its own set of commands, but certain commands may be the same in all of them.

A/UX editor commands are usually single characters that stand for a function. It may be somewhat difficult to remember the abbreviation for each command, but for the most part the command names are **mnemonic.** That is, the abbreviation for each command corresponds to the function of the command. For example, `d` stands for `delete`, `w` stands for `write`, and `q` stands for `quit`.

Entering text

When you open a file, the editor copies the file into the **editing buffer.** The editing buffer is a temporary workspace similar to a blank sheet of paper. When you create a file, you insert text into the buffer. When you modify a file, you make the changes to a copy of the file in the buffer.

You can modify the text in the buffer, insert new text, delete text, move blocks of text to new locations, substitute one string of text for a second at every occurrence of the first, and so on. Remember that everything you do to the buffer contents is temporary until you write the contents back to the file.

Changing the default editor

As shipped, TextEditor is the default editor for A/UX. This means that when you double-click the icon of a text-only file, TextEditor opens the file.

The default editor can be different for different accounts. For example, you can set the text editor for the root account to `vi`, and leave TextEditor as the default editor for user accounts.

Each time you log in to a user account or the root account, the system creates your working environment by reading a set of environment variables. These environment variables are stored in the `.login` file if you use the C shell (the default shell for user accounts) or in `.profile` if you use the Korn or the Bourne shell. You change the default editor by setting an environment variable in one of these files.

Changing the default editor for a user account

Follow these steps to change the default editor for a user account:

1 **Log in to the root account.**

   You must be logged in to the root account to edit the `.login` or `.profile` files.

   If you use the default shell for user accounts, your shell is the C shell; skip to step 3.

2 **Find out which shell your account uses by looking in** `/etc/passwd`.

   To open a CommandShell window, choose CommandShell from the Apple ( ) menu and New from the File menu. Enter

   ```
   cat /etc/passwd
   ```

   Find the line beginning with your account name. The last characters on this line specify the shell in use by your account: `csh` for the C shell, `sh` for the Bourne shell, and `ksh` for the Korn shell.

3 If your account uses the C shell, open the `.login` **file in your account's home directory.**

To make the Finder active, choose Finder from the Apple menu. Open your account's home directory by double-clicking its folder. Open `.login` by double-clicking it. Skip to step 5.

4 If your account uses the Bourne or Korn shell, open `.profile` **in your account's home directory.**

To make the Finder active, choose Finder from the Apple menu. Open your account's home directory by double-clicking its folder. Open `.profile` by double-clicking it.

5 Under the line beginning with `setenv TERM` **add a new line as follows:**

`setenv FINDER_EDITOR` *editorpathname*

Replace the word in italics with the full pathname of the editor you wish to make the default editor. For example, to set `vi` as your default editor, type

`setenv FINDER_EDITOR /usr/bin/vi`

Or to set TextEditor as your default editor, type

`setenv FINDER_EDITOR /mac/bin/TextEditor`

6 Save the change and close the file.

7 To effect the change, log out and log back in to the account.

Changing the default editor for the root account

Follow these steps to change the default editor for the root account:

1 Log in to the root account.

You must be logged in to the root account to edit the `.login` or `.profile` files.

If you use the default shell for the root account, your shell is the Bourne shell; skip to step 4.

2 **Find out which shell the root account uses by looking in** `/etc/passwd`.

To open a CommandShell window, choose CommandShell from the Apple menu and choose New from the File menu. Enter

```
cat /etc/passwd
```

Find the line near the top of the file beginning with `root`. The last characters on this line specify the shell in use by the root account: `csh` for the C shell, `sh` for the Bourne shell, and `ksh` for the Korn shell.

3 **If the C shell is in use, open the** `.login` **file for the root directory.**

Double-click the / disk to open it. Double-click `.login` to open this file.

Skip to step 5.

4 **If the Bourne or Korn shell is in use, open** `.profile` **for the root directory.**

Double-click the / disk to open it. Double-click `.profile` to open this file.

5 **Under the line beginning with** `setenv TERM` **add a new line as follows:**

```
setenv FINDER_EDITOR editorpathname
```

Type the line, and replace the word in italics with the full pathname of the editor you wish to make the default editor. For example, to set `vi` as your default editor, type

```
setenv FINDER_EDITOR /usr/bin/vi
```

Or to set TextEditor as your default editor, type

```
setenv FINDER_EDITOR /mac/bin/TextEditor
```

6 **Save the change and close the file.**

7 **To effect the change, log out and log back in to the root account.**

2 Using TextEditor

This chapter provides a detailed description of the features and capabilities of TextEditor, the only mouse-based text-editing program included with A/UX. TextEditor is the default A/UX editor and is compatible with all text-only files. In other words, TextEditor can open and edit files created by `vi, ed, ex,` and `sed`.

What is TextEditor?

TextEditor is the default editor for A/UX. It allows you to create and edit text files using the mouse and menu commands in the traditional Macintosh manner. If you are just learning to use A/UX and the different text-editing tools it provides, TextEditor is the best place to start.

TextEditor creates a file that contains only the text characters that you type on your keyboard (including tab and return characters). This is called a **text-only file** (or an **ASCII text file**). Any file you create with word-processing or desktop-publishing software, unless you save it as "text only," contains many formatting commands. Although these formatting commands are not visible on your screen, they confuse other programs that try to use your file. Any text editor or word processor, including vi, ed, ex, and sed, can read text-only files. Moreover, you can use the editor to write programs or shell scripts, which cannot contain hidden formatting characters.

To preserve the text-only nature of TextEditor's text files and still allow you certain formatting choices, TextEditor gives you the option of saving the file's formatting information. It saves formatting information in a separate file called a **resource fork** and saves the text characters in a text-only file. These formatting choices include adjusting the tab settings and choosing fonts, among others. These features are described later in this chapter.

Because TextEditor reads and writes text-only files, you can use it to work with files created with any other text-only text editor. Thus, you can edit MacWrite® files saved with the Text-Only option. When you choose the Open command, TextEditor's standard File dialog box displays a list of all text files in the folder (or the "current directory," in UNIX parlance), regardless of what program was used to create the file.

The resource fork has the same filename as the text file, but it is preceded by a percent sign (%).

Like most UNIX systems, A/UX 3.0 also comes with a commonly used UNIX text editor called vi. See Chapter 3, "Using the vi Screen Editor," for more information. If you are an experienced vi user, you can continue to use vi to read and write text files from the CommandShell command line. Naturally, TextEditor can read and edit text files created with vi, and vice versa.

Creating a new file

The following sections teach you how to start TextEditor and create a new file.

Starting TextEditor

You start TextEditor from a CommandShell window by following these steps:

1 **Choose CommandShell from the Apple menu.**

The Apple menu is located at the far left of the menu bar.

2 **Type** `TextEditor` **at the prompt.**

This command is case-sensitive, so make sure you type the correct capital letters.

3 **Press RETURN.**

The TextEditor application starts, and an untitled window appears.

Entering text

Of all the interactive editors included with A/UX, TextEditor is the only one that takes full advantage of the mouse. You can use the mouse to move the cursor anywhere on the screen. You can highlight text for deletion and insertion, or place the cursor in a specific location in a file.

In TextEditor, you insert text just as you do in most common word-processing programs. Text you type at the keyboard appears on the screen at the location of the blinking text-insertion point, or **I-beam.** By default, the I-beam appears at the very top of the screen when you first open a new or existing file.

Try typing a few lines of text. Notice that the text does not automatically move to the next line when you type past the right edge of the screen. Use the mouse to move the cursor to the beginning of a word that is near the right edge of the screen; then press RETURN. You may also move the insertion point to the left by pressing and holding the LEFT-ARROW key.

Cutting and pasting text

One of TextEditor's most useful features is the ability to cut, paste, and move text. To select a specific range of text, you must **highlight** it. To highlight or select text, follow these steps:

1 **Place the I-beam before the first letter of the text you wish to select.**

2 **Hold down the mouse button and move the mouse so that the I-beam moves to the end of the desired text.**

3 **Release the mouse button.**

The highlighted text appears as light letters on a dark background, rather than dark on light.
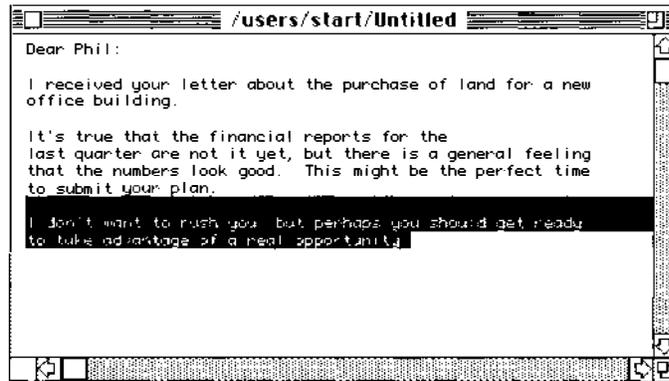
You can perform various operations on a block of selected text. If you perform a Cut or Copy operation, the selected text is sent to the **Clipboard.** The Clipboard is a temporary buffer in which TextEditor stores text. You can view the contents of the Clipboard by choosing Show Clipboard from the Edit menu.

Using the Cut command

The Cut command copies any selected text to the Clipboard and removes the text from the window. The text is stored in the Clipboard until it's replaced by text sent there as a result of another Cut or Copy command. The Command-key equivalent is COMMAND-X. Since the text is removed from the screen, the Cut command is most commonly used to move blocks of text within a file.

Using the Copy command

The Copy command copies any selected text to the Clipboard but does not remove it from the window. The text is stored in the Clipboard until it's replaced by text sent there as a result of another Cut or Copy command. The Command-key equivalent is COMMAND-C. Since the text isn't removed from the screen, the Copy command is most commonly used to move blocks of text between different files.

Using the Paste command

The Paste command inserts the contents of the Clipboard into the window at the insertion point. The Command-key equivalent is COMMAND-V. You may paste the contents of the Clipboard as many times as you wish, but remember that the contents of the Clipboard change the next time you give a Cut or Copy command.

Saving a TextEditor file

Four commands allow you to save a TextEditor file. The first is Save. If the file is already named, simply choose Save from the File menu or press COMMAND-S to save the current document. If you haven't given the file a name, TextEditor displays a dialog box prompting you to type a name for the new file.

Using additional commands to save a TextEditor file

In addition to Save, TextEditor offers three commands that allow you to save your document. They are Save As, Save a Copy, and Revert to Saved.

The Save As command

The Save As command displays a dialog box that allows you to make a copy of the currently active file, which you must then save under a different name. This action saves the current contents of the window under the new filename and allows you to continue editing the new file. The old file is closed without saving, under its original name. You generally use this command when you are working on a file that already has a name. This way, you can save the edited file under a new name, leaving the original intact.

The Save a Copy command

The Save a Copy command displays a dialog box that allows you to save the current state of the active window to a new file with the name Copy Of *Filename*. You can then continue editing the old file.

The Revert to Saved command

Choose the Revert to Saved command if you want to throw away any changes you have made since you last saved the file in the active window. This command is dimmed if the file has not been modified since you last saved it.

Editing an existing file

You can use TextEditor to open text files created with `vi`, `ed`, and `ex`. Later chapters describe the functions of the different editors.

Opening a text file

Suppose you want to edit a text file that already exists. If TextEditor is running, you can open the file for editing by choosing Open from the File menu. This command allows you to open any text file, regardless of what application was used to create the file.

If you are using the Finder operating system software and if TextEditor is your system's default text editor, you need only double-click the icon of the file to be edited (or click it once to select it, and then choose Open from the File menu). TextEditor starts running, and the file you have clicked is opened for editing.

◆ **Note** To open a file, you need to have read permission to it. If you open a file for which you have read but not write permission, a pencil with a slash over it appears at the bottom left of the window, as shown in the following figure. You can read the document but you cannot change it. File access permissions are explained in *A/UX Essentials.* ◆


```
▉▉▉▉▉▉▉▉▉▉ /•cshrc ▉▉▉▉▉▉▉▉▉▉
# .cshrc
# sample of commands commonly found in .cshrc
# The .cshrc file is run everytime you invoke Csh
# We have found the following settings to be useful
# Feel free to make modifications and add your own

    # set the login shell prompt
if ($?prompt) then               # interactive shell
   # set prompt string
   switch ("$prompt")
     case '#':
     case '%':
        set prompt="`hostname`.$LOGNAME \! $prompt"
   endsw
else
   exit
endif

    # aliases: use a personal alias in place of the command it re
```

Write permission denied

Using TextEditor's search-and-replace commands

One of the most important time-saving features of an interactive text-editing application is its ability to search through a document for a specific word, phrase, or string of characters and to change it automatically. With other interactive editors such as `vi`, you must remember certain commands and options to perform these search-and-replace tasks. TextEditor has built-in menu commands that make finding and changing text quick and easy.

Using the Find command

This command displays the Find dialog box and searches for the string that you type in the "Find what string?" text field. By default, the editor searches forward from the currently selected text in the active window (and does not wrap around). The Command-key equivalent is COMMAND-F. The Find dialog box has five options.

```
┌─────────────────────────────────────────────┐
│  Find what string?                          │
│ ┌─────────────────────────────────────────┐ │
│ │ GOOD                                    │ │
│ └─────────────────────────────────────────┘ │
│ ○ Literal              ⊠ Case Sensitive     │
│ ◉ Entire Word          ☐ Search Backwards   │
│                        ☐ Wrap-Around Search │
│ ┌──────────┐              ┌──────────┐       │
│ │   Find   │              │  Cancel  │       │
│ └──────────┘              └──────────┘       │
└─────────────────────────────────────────────┘
```
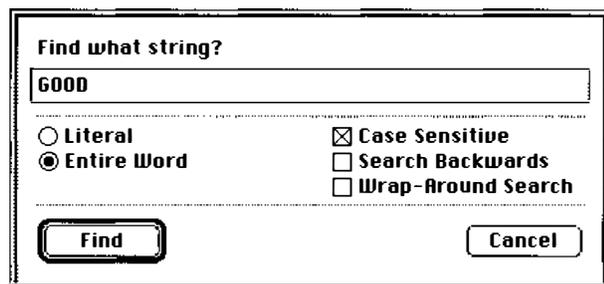
Literal: Click the Literal button to search for a string of characters even when they are embedded in another word. Thus, if you click Literal and search for the word *it,* the system finds *bit, split, flit,* and so on.

Entire Word: Click the Entire Word button to search for whole words only.

Case Sensitive: Click the Case Sensitive button if you want to find only words that exactly match the case (capitalization and lowercase) of the word in the search string.

Search Backwards: Click Search Backwards to search backward from the current selection to the beginning of the file. (Normally, the search moves forward and stops at the end of the file.)

Wrap-Around Search: Click Wrap-Around Search to search forward from the location of the cursor to the end of the file, then start again from the beginning to the starting cursor position.

Cancel: Removes the dialog box. No further action is taken.

◆ **Hints on using Find** You can reverse the direction of a current search operation by pressing SHIFT as you click the Find or OK buttons. For example, if you are in the middle of a file and you want to find a string that occurs earlier in the document, hold down the SHIFT key as you click Find. The search then proceeds backward through the first part of the file. The direction is changed for the current search only. ◆

Using the Find Same command

This command repeats the last Find operation on the active window. The Command-key equivalent is COMMAND-G.

Using the Find Selection command

This command finds the next occurrence of the search string in the active window. The Command-key equivalent is COMMAND-H.

Replacing text

After you have found a text string, you may want to replace it with another string of text. For instance, you would use the Replace command if you wanted to replace a specific occurrence of the word *GOOD* in a file with the word *excellent*.
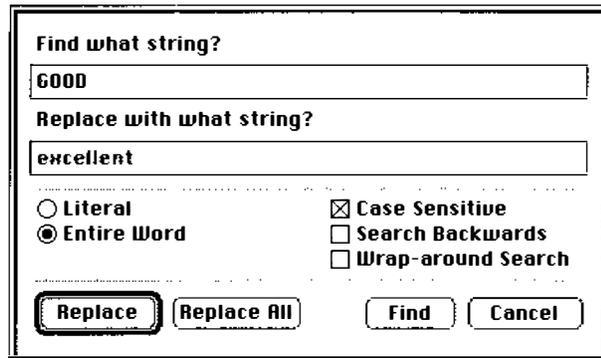
```
Find what string?
┌────────────────────────────────────────┐
│GOOD                                     │
└────────────────────────────────────────┘
Replace with what string?
┌────────────────────────────────────────┐
│excellent                                │
└────────────────────────────────────────┘

○ Literal            ⊠ Case Sensitive
◉ Entire Word        ☐ Search Backwards
                     ☐ Wrap-around Search

[ Replace ]  [Replace All]   [ Find ]  [ Cancel ]
```

Using the Replace command

This command searches for a specified string throughout a file and replaces that string with a different string. The Command-key equivalent is COMMAND-R. To replace a text string, follow these steps:

1 **In the field labeled "Find what string?" type the text string that will take the place of the search string.**

In this example, the word to find is *GOOD*.

2 **Type the text string you wish to replace it with in the field labeled "Replace with what string?"**

In this example, *excellent* replaces *GOOD*.

3 Click Find.

TextEditor begins searching forward in the file from the location of the I-beam. Remember, if Search Backwards has been clicked, the search occurs in reverse order. Once TextEditor finds the search string you designated earlier, (the word *GOOD*), the search stops, and the string is highlighted in the file.

4 Click Replace (or press RETURN).

In this example, the word *excellent* replaces the word *GOOD,* and the next occurrence of *GOOD* is highlighted after a few moments, allowing you to repeat the operation.

If you want to skip this occurrence of the string without changing it, click the Find button. TextEditor leaves the highlighted word unchanged and searches for the next occurrence.

The Replace dialog box offers this additional option:

Replace All: Changes all occurrences of the specified string automatically.

▲ **Warning** Before using this command be sure to click the Entire Word button to ensure that you don't change embedded instances of the text string. ▲

Using the Display Selection command

This command causes the currently selected text in the active window to scroll into view.

Using the Replace Same command

This command repeats the last Replace operation. The Command-key equivalent is COMMAND-T.

Formatting and other features

If you use any formatting features, you can save the formatting information by selecting Save Formatting Information in the Save As dialog box. This creates and saves an extra file, called a resource fork, that contains the formatting information. This feature allows you to retain a text-only copy of the file that is compatible with any A/UX editor.

Changing fonts

TextEditor comes with six fonts: Chicago, Courier, Geneva, Helvetica®, Monaco, and the default font, Times®. You may use the Font D/A Mover application to move fonts to your A/UX system. To learn more about adding fonts, see "Customizing Your Work Environment," in *A/UX Essentials*.

You may change the font of a character, a word, or an entire file. To change fonts, follow these directions:

1 **Choose Format from the Edit menu.**

The Format dialog box appears.
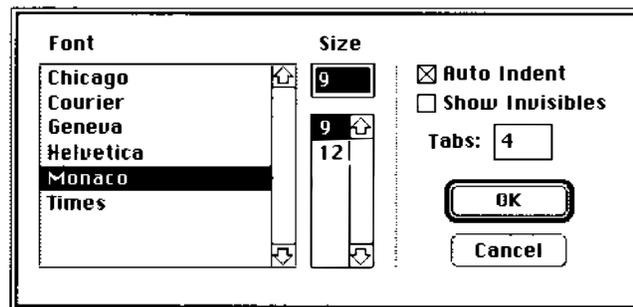
2 **Select the desired font and size.**

3 **Click OK.**

Selecting tab settings

Tabs are automatic stops used to set margins in a file. You insert tabs by pressing the TAB key. Tabs are primarily used for formatting tables. To change tab settings, follow these directions:

1 Choose Format from the Edit menu.

The Format dialog box appears.

2 To change the tab setting, type the desired number of spaces in the Tabs field.

3 Click OK.

You may need to experiment with different settings until you find one that fits your needs.

Automatically aligning text

The Auto Indent option aligns a selected block of text or a line with the previous line. To turn Auto Indent on, follow these steps:

1 Choose Format in the Edit menu.

The Format dialog box appears.

2 Click Auto Indent.

An X shows in the checkbox. You can click it again to remove the X, thus turning Auto Indent off.

3 Click OK.

Showing invisible characters

Space, tab, return, and control characters don't appear on your screen unless you click Show Invisibles. Table 2-1 shows the corresponding screen symbol for each of the invisible characters.

Table 2-1 Invisible characters

Character	Screen symbol
tab	Δ
space	◊
return	¬
all other control characters	¿

To make the invisible characters appear on the screen, follow these steps:

1 **Choose Format from the Edit menu.**

The Format dialog box appears.

2 **Click Show Invisibles.**

An X shows in the checkbox. You can click it again to remove the X, thus hiding the invisible characters.

3 **Click OK.**

Shifting text left

The Shift Left command moves any selected text in a TextEditor window one tab stop to the left. The Command-key equivalent is COMMAND-[.

Shifting text right

The Shift Right command moves any selected text in a TextEditor window one tab stop to the right. The Command-key equivalent is COMMAND-].

◆ **Note** If you hold down the S<small>HIFT</small> key while choosing Shift Left or Shift Right, the selected text shifts by one space rather than by one tab stop. ◆

Arranging multiple windows

If you are editing several documents at once, you have a choice of stacking or tiling the windows. The commands discussed in this section are found in the Window menu.

Tiling windows

When you choose Tile Windows from the Window menu, TextEditor displays any windows on the desktop.
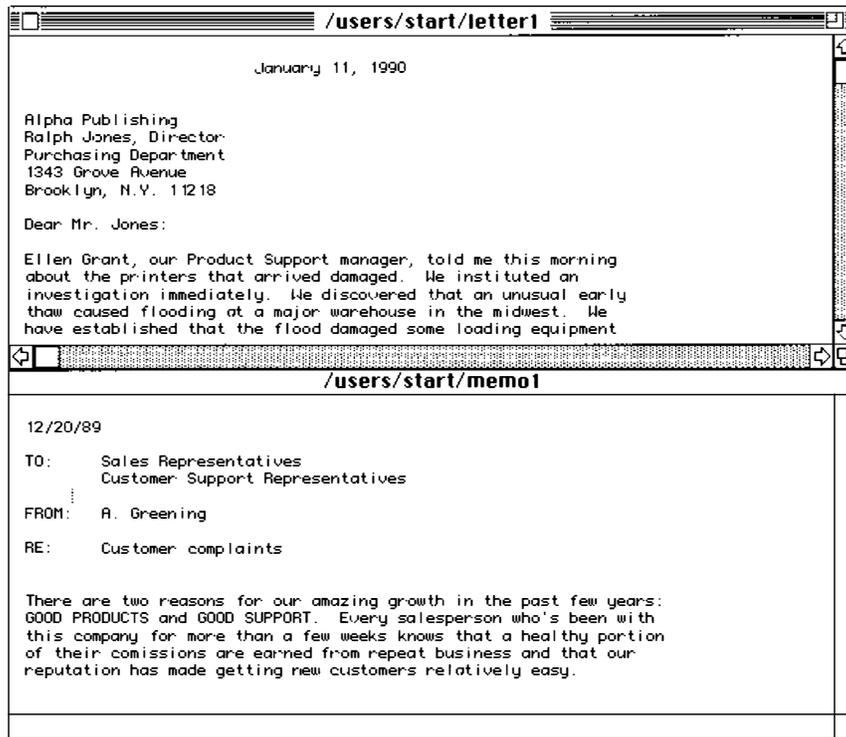
Stacking windows

When you choose Stack Windows from the Window menu, TextEditor stacks any windows on the desktop, as shown the figure below.
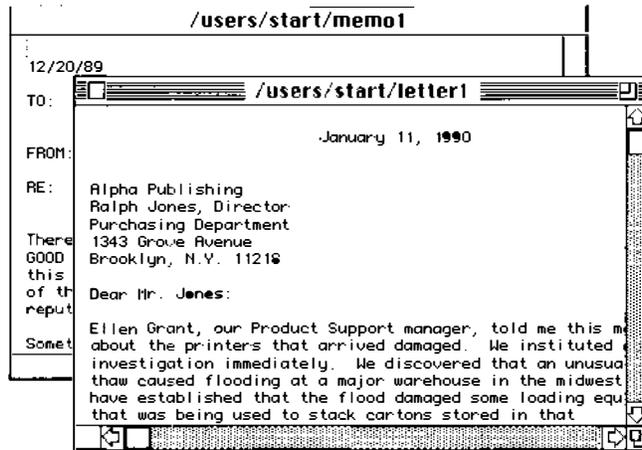
△ **Important** When a window becomes too small, the text in the window disappears and is lost. You have to reopen the file in a larger window. ⁄

Marking a place in the file

If you need to return periodically to specific locations in a long or complex document, you can place invisible named markers at those points. You can then jump instantly to any of the marked places in the document. This saves you the time and trouble of having to scroll through the document to find those places.

1 **Highlight the word or words you wish to mark.**

2 **Choose Mark from the Mark menu.**

The dialog box appears.

```
┌─────────────────────────────────────────────────────────────┐
│  Mark the selection with what name?                         │
│ ┌─────────────────────────────────────────────────────────┐ │
│ │ We                                                       │ │
│ └─────────────────────────────────────────────────────────┘ │
│  ┌─────────────────┐            ┌─────────────────┐         │
│  │       OK        │            │     Cancel      │         │
│  └─────────────────┘            └─────────────────┘         │
└─────────────────────────────────────────────────────────────┘
```
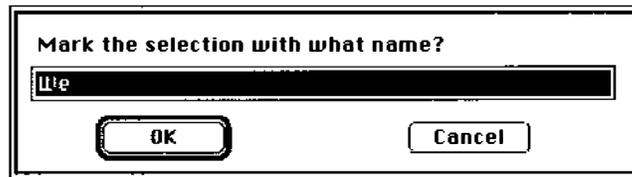
You can use the highlighted word or words already in the text box as the name of the marker. If you want to give the marker a different name, type it in the text box.

3 Click OK.

You have named the first marker *We*.

The named marker appears in the Mark menu. Each time you choose the marker, you return to that location in the file.

Removing markers from text

To remove markers, follow these steps:

1 Choose Unmark in the Mark menu.

The dialog box appears.

```
┌─────────────────────────────────────────────────┐
│  Delete which markers?                          │
│ ┌──────────────────────────────────────────┐ ▲  │
│ │ We                                       │    │
│ │ 2                                        │    │
│ │                                          │    │
│ │                                          │    │
│ │                                          │ ▼  │
│ └──────────────────────────────────────────┘    │
│  ┌─────────────┐        ┌─────────────┐         │
│  │   Delete    │        │   Cancel    │         │
│  └─────────────┘        └─────────────┘         │
└─────────────────────────────────────────────────┘
```
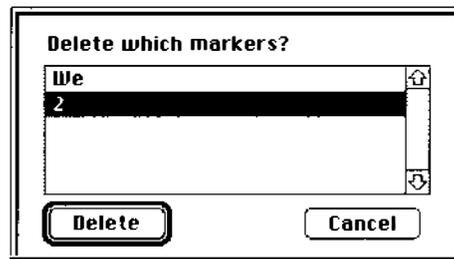
2 Highlight the name of the marker you wish to remove.

3 Click Delete.

The marker is removed from the Mark menu.

Printing

TextEditor allows you to print either an entire document or a selection from a document. Your computer needs to be properly connected to a printer or to a network with a printer. You also need to have a printer selected in the Chooser (in the Apple menu). For information on using the Chooser to select a printer, see *A/UX Essentials*. For information on connecting to a network to use its printer, see *A/UX Networking Essentials*.

Printing an entire document

To print an entire document, follow these steps:

1 **Open the document in TextEditor.**

2 **Choose Print Window from the File menu.**

You see the Print dialog box.

```
LaserWriter "Cedar Point"                          7.0▶19   [ Print ]

Copies: [1]          Pages: ⦿ All  ○ From: [    ] To: [    ]   [ Cancel ]

Cover Page:    ⦿ No ○ First Page  ○ Last Page            [ Help ]

Paper Source: ⦿ Paper Cassette  ○ Manual Feed

Print:         ⦿ Black & White   ○ Color/Grayscale

Destination:   ⦿ Printer         ○ PostScript® File

☐ Print Back to Front
```
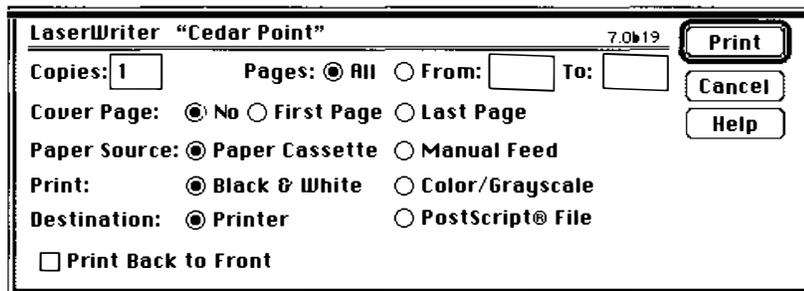
3 **Click OK to begin printing.**

Printing a portion of a document

To print a portion of a document, do the following:

1 **Highlight the text block you want to print.**

2 **Choose Print Selection in the File menu.**

The Print dialog box appears.

3 **Click OK to start printing.**

Quitting TextEditor

You may want to quit TextEditor to perform other tasks or to take a break. To quit TextEditor, simply choose Quit from the File menu or type COMMAND-Q.

If you have made changes to a file, you see a Save Before Quitting dialog box, which asks whether you want to save those changes.

```
  ⚠   Save changes to /users/start/letter1

  [ Don't Save ]      [ Save ]      [ Cancel ]
```

Click Yes to save, click No to quit without saving, or click Cancel to cancel the Quit command.

To save the document as a text-only file, click Save Text Only. To save a resource fork with formatting information (in addition to the text file), click Save Format Information. Remember, since TextEditor automatically saves all formatting information in a resource fork file, you don't need to save the file as text only to make the file compatible with the other editors discussed in this book.

3 Using the `vi` Screen Editor

This chapter provides a detailed description of the features and capabilities of the UNIX-standard screen editor `vi`.

What is `vi`?

Of all the text-editing programs A/UX has to offer, `vi` is perhaps the most powerful. If you are new to A/UX and the UNIX operating system, `vi` may appear somewhat difficult to understand. Veteran UNIX users can attest to the fact that once you get used to `vi`, you'll find that its many options make it a useful and powerful tool for all your editing needs.

The name `vi` (pronounced *vee-eye*) is derived from the word *visual*. It is a visual editor in the sense that it uses the full screen as a window into the file you edit. When you make a change, `vi` immediately displays the change on your screen.

In addition to `vi` two other commands, `view` and `vedit`, use the same interactive visual style as `vi`.

- `view` is similar to `vi` but protects you from making unintended changes by setting read-only permission on the file.

- `vedit` is identical to `vi` except that you see an `INPUT MODE` message when you are entering text, and `vedit` reports the number of changes you make with global substitutions (when the number of changes is greater than one). The `vedit` command is intended for beginning users, but most users prefer to use `vi` and `view` exclusively.

The relationship between `vi` and `ex`

Although `vi` and `ex` can be used as separate text-editing programs, they are actually two aspects of one program. To understand their relationship, think of `vi` as a visual interface that works by displaying the contents of a file as you edit. Next, think of `ex` as a set of commands that `vi` accesses to accomplish certain tasks that it otherwise could not accomplish on its own. In most cases, `vi` and `ex` work best as a team.

- `vi` is a screen editor. You type commands to add or change text anywhere on the screen, and the screen changes immediately to show the changes. Most of the time, you don't need to know the line numbers of the lines you wish to work on. In this book, the terms "`vi`" and "visual mode" both refer to the `vi` screen editor. TextEditor is another example of a screen editor.

- ex is a line editor. A line editor works by specifying a set of lines on which to operate (for example, add text after this line, or change these ten lines). You issue commands to add or change text in response to a command prompt, and you cannot always see the results of changes right away. In most cases, you'll need to know the line numbers of the lines you wish to modify or otherwise operate on. In this book, the terms "ex" and "line mode" (which is usually accessed within vi) both refer to the ex line editor. Another example of a line editor is ed.

Starting vi

There are two methods for starting vi. The standard method is to open a CommandShell window, type vi *filename,* and then press RETURN. This is the quickest way to start vi, but vi cannot find the file you wish to open unless you know the filename. If you type the wrong filename or make a typing error, vi creates a new file with that name. For example, if you want to open an existing file named bernard and you accidentally type bernardo, vi opens an empty file named bernardo. To avoid this problem, you can use Commando to start vi and open a new or existing file. Commando allows you to choose the file you wish to open directly from a search box.

Starting vi and opening an existing file

If you want to open an existing file, follow these steps:

1 **Log in to A/UX.**

2 **Choose CommandShell from the Apple menu.**

The Apple menu is located at the far right of the menu bar.

3 Type `vi` at the shell prompt and press COMMAND-K.

The `vi` Commando dialog box appears.

```
┌─vi Options ─────────────────────────────────────────────────┐
│ ┌─Operation ──────────────────┐    File to recover and edit:   │
│ │ ● Create new file to edit    │   ┌──────────────────────────┐│
│ │ ○ Edit existing file         │   └──────────────────────────┘│
│ │ ○ Recover and edit lost file │    Specify tag:                │
│ │ ○ Edit file containing 'C' tag│  ┌──────────────────────────┐│
│ └──────────────────────────────┘  └──────────────────────────┘│
│   ┌──────────────────────────┐     ☐ Lisp mode                 │
│   │  Enter name of new file   │     ☐ Read only mode            │
│   └──────────────────────────┘     ☐ Encryption mode            │
│   ┌──────────────────────────┐    ┌──────────────┐┌───────────┐│
│   │   Choose file(s) to edit  │    │ More options ││Output & Error││
│   └──────────────────────────┘    └──────────────┘└───────────┘│
│ ┌─Command line ─────────────────────────────────────────────┐ │
│ │ vi                                                         │ │
│ └────────────────────────────────────────────────────────────┘ │
│ ┌─Help ────────────────────────────────────┐  ┌─────────────┐ │
│ │ vi is a screen-oriented (visual) display text editor. │  │   Cancel    │ │
│ │                                           │  └─────────────┘ │
│ └───────────────────────────────────────────┘  ┌─────────────┐ │
│                                                 │     vi      │ │
│                                                 └─────────────┘ │
└──────────────────────────────────────────────────────────────┘
```
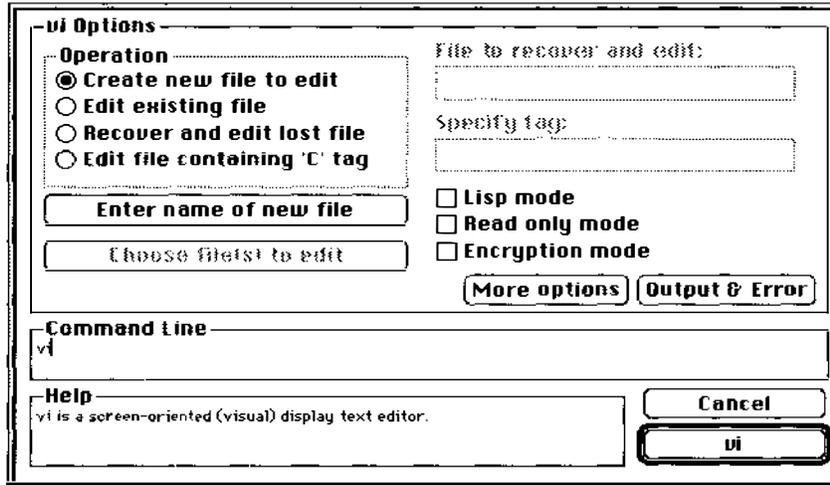
4 Click "Edit existing file."

You will need to choose the file you wish to open.

5 Click the "Choose file(s) to edit" button.

A standard Macintosh search box, like the one shown on the next page, appears.
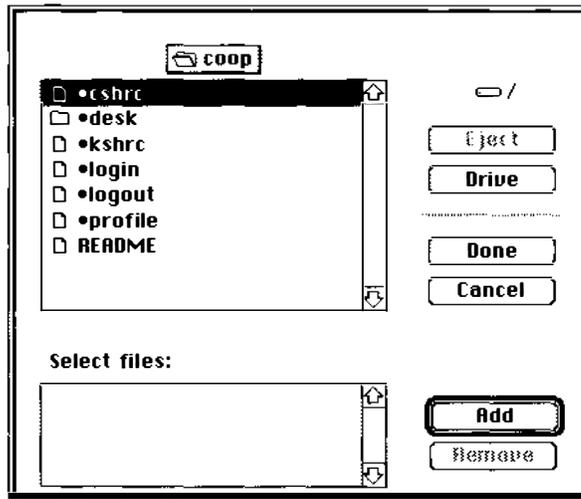
6 **Double-click the file you wish to open.**

To open more than one file, for each file select the file and click the Add button.

You return to the main vi Commando dialog box.

7 **Click "Done."**

8 **Click the "vi" button.**

You return to the CommandShell window.

9 **Press RETURN.**

The vi application starts, and the file or files you have selected are opened and ready for editing.

You can also use Commando to create a new file. To do this, open the vi
Commando dialog box, click the "Enter name of new file" button, type the name of the file, then click OK. Click the "vi" button and press RETURN to open the new file.

Syntax and initialization

The following sections describe the command syntax and initialization procedures for the `vi` editor.

Command syntax for the `vi` editor

The command syntax for `vi` is as follows:

`vi` [+*command*] [-l] [-r [*filename*]] [-R] [-t [*tag*]] [-w*n*] [-x] [*filename* ...]

You can also use the `view` and `vedit` commands with the same flag options.

The options are as follows:

+*command*
: Move to the line specified by *command*, where *command* is either a regular expression (see "Regular Expressions and Searching," later in this chapter) or a line number (for example, `+100` starts editing at line 100). If you omit *command*, `vi` moves the cursor to the last line of the first file.

-l
: Set the `showmatch` and `lisp` options for editing LISP programs. These are described in "Setting Options," later in this chapter.

-r [*filename*]
: Recover a file after an editor or system crash. If you omit *filename*, `vi` lists the saved files.

-R
: Set the `readonly` option, making it impossible to write the file with the `write` command.

-t [*tag*]
: Start editing the file at *tag* (usually a spot marked with the `ctags` program). Equivalent to an initial `tag` command. This is described in Chapter 4, "Using the `ex` Line Editor."

-w*n*
: Set the window size to *n* lines.

-x
: Prompt for a *key* to encrypt and decrypt the file (see `crypt`(1) in *A/UX Command Reference*). The file should already be encrypted using the same key.

filename
: The file(s) to edit.

Initialization procedures of the `vi` editor

Upon startup `vi` sets up your editing environment by taking the following steps:

- reads the `TERM` variable to find out what terminal you're using
- sets any options you've specified in the `EXINIT` environment variable, usually set in the `.profile` (or `.login`) file in your home directory
- sets any options you've specified in the `.exrc` file in the current directory or your home directory

You can set the same options with either the `EXINIT` environment variable or `.exrc` files. The options are described in "Setting Options," later in this chapter.

Opening a file

You can display a file for **read-only viewing,** or for editing and making changes. Read-only viewing means that you can read the file but cannot make any changes to it. File permissions are discussed in *A/UX Essentials* and in *A/UX Networking Essentials*.

Read-only viewing

If you want to look at a file while protecting it from unintended changes, start the viewing program from the shell by typing

`view` *filename*

instead of `vi` *filename*. The `view` command protects the file from accidental changes. When you enter a file using `view`, you can use all of `vi`'s commands, but you can make changes to the file only by typing the colon character (`:`) to move to line mode and using the command

`:w!`

You can then exit the file by typing the colon character (`:`) to move to `ex` line mode and typing

`q`

If you exit `view` using the `vi` command

`ZZ`

you exit the file without making any permanent changes. If you try to exit `view` using commands that write changes to a file before exiting `vi` (for example, `:wq`), you get the following error message:

`File is read only`

Type

`q!`

to quit the application.

Opening a file for editing

To create and open a new file (or open an existing file) in `vi`, type

`vi` *filename*

where *filename* is the name of the file you're creating (or opening).

When you use `vi` to create a new file, `vi` opens a temporary storage file called a **buffer.**

When you edit an existing file, `vi` places a copy of that file in the buffer. Changes are made only in this temporary copy in the buffer. The editor does not change the actual contents of the file until you save your changes. See "Saving Files and Quitting," later in this chapter, for more information.

⚠ **Important** You should periodically save your changes to the file to prevent losing material if the system crashes or an error occurs. ⚟

The different modes of `vi`

The `vi` editor has three modes.

- When you first open a file, you are in `vi` command mode. You use `vi` command mode to perform basic editing tasks such as deleting and moving text, as well as file maintenance tasks such as saving and quitting. In `vi` command mode, the screen does not display the commands you type on the keyboard. In this mode, `vi` assumes anything you type is a command and tries to run it. For example, suppose you've created and opened a new, empty file and you type the letter `j`. Because `j` is a cursor-movement command (discussed later in this chapter), `vi` tries to move the cursor accordingly. Because there is no text in the file, there is no place to move, and `vi` signals that it cannot comply with the command.

- To insert text, you must enter insert mode. This is the mode you use to do most of your basic text input. In insert mode, `vi` assumes anything you type is text (rather than commands). To enter insert mode, type `i`. After that, anything you type appears on the screen. The `vi` editor places what you type in the buffer. Your text is not written to the file until you return to command mode and save the file (see "Saving Files and Quitting `vi`"). Other `vi` commands for inserting text include `a` (append), `o`, and `O` (open line). There are also several commands, such as `c` (change) and `s` (substitute), that insert text. For a complete list of `vi`'s insertion commands, see "Inserting Text," later in this chapter. You always leave insert mode by pressing the ESCAPE key.

- You can use the `ex` editor commands by entering `ex` command mode, sometimes known as line mode. See "Switching to `ex` Command Mode," next.

Switching to `ex` command mode

You use `ex` command mode to issue `ex` commands from within `vi`. To use `ex` commands from within `vi`, first press the ESCAPE key (if necessary) to enter `vi` command mode, then type the colon character (`:`). The `ex` command line appears at the bottom of the screen, ready to accept any `ex` command. You use `ex` commands from within `vi` to perform global changes, searches, and other operations that involve more than one line in the file.

You can also switch to line mode for a series of `ex` commands (or a multiple-line `ex` command) by typing

`Q`

in `vi` command mode. To return to `vi` when you have entered line mode this way, type

`vi`

at the colon prompt on the bottom line of your screen.

◆ **Note** You can also switch to the `ex` command line to run `ex` search and global commands by typing the slash character (`/`) or the question mark (`?`). ◆

For complete information on using `ex` commands, see Chapter 4, "Using the `ex` Line Editor."

Using special keys in `vi`

The following keys have special meaning in A/UX.

ESCAPE	The ESCAPE key ends all text insertion in `vi` and returns you to command mode.
RETURN	The RETURN key ends all commands given on the `ex` command line. See "Switching to `ex` Command Mode," earlier in this chapter.
The interrupt key sequence	The interrupt key sequence (CONTROL-C in A/UX) sends an interrupt signal to the editor. It gives you a forceful way of stopping `vi` from executing a command it has already started.

The `vi` editor occasionally shows your commands on the last line of the screen. If the cursor is on the first position of this last line, `vi` is working on something (such as finding a new position in the file after a search or reformatting the buffer). When this happens, you can stop `vi` by sending an interrupt.

Displaying text and moving within a file

To move the cursor within `vi`, you use either arrow keys or their equivalents, or motion commands. Users who have come to rely on the mouse to move the cursor might want to take a few minutes to practice using the keyboard and motion commands.

▲ **Warning** Using the mouse to resize a window may result in lost data. ▲

Using arrow keys to move within a file

The arrow keys on your keyboard move the cursor in `vi`. The `h, j, k,` and `l` commands also move the cursor.

`h` or ←	Move the cursor left a space. (The system *erase* key, usually DELETE, also works.)
`l` or →	Move the cursor right one space. (Space bar also works.)
`k` or ↑	Move the cursor up a line (in the same column).
`j` or ↓	Move the cursor down a line (in the same column).

You use these keys in command mode, and you can precede them with a number indicating how many spaces you want to move in the direction you want. For example, `5h` moves the cursor left five spaces.

Using motion commands to move within a file

Motion commands are either mnemonic, single-character commands or symbols that move the cursor in a file without affecting the file's contents in any way. With the exception of the `G` command, motion commands operate relative to the current cursor position. You can combine motion commands with a number (to indicate how many times the command runs) or with an operator, such as `d` for delete (to indicate how far the operation extends).

If preceded by a number *n,* a motion command moves *n* motions (for example, *n* spaces or *n* lines) in that direction. The syntax is then

[*n*] *motion*

For example, to move the cursor three words forward, type

3w

Preceding these commands with an operator, such as d (delete) or c (change), indicates how far the operation of deleting or changing text extends. In this case, the syntax is

[*n*] [*opr*] *motion*

For example, to replace two words of text, type 2cw. See "Combining Operators and Motions," later in this chapter, for more details.

Table 3-1 Motion commands in vi

Command	Description
[*n*]-	Move the cursor to the beginning of the preceding line. Scroll if necessary.
[*n*]+	Move the cursor to the beginning of the next line. Scroll if necessary.
[*n*]$	Move the cursor to the end of the current line. Preceded by a number it means "move to the end of the line *n* lines forward in the file."
^	(Caret.) Move the cursor to the beginning of the first word on the line.
0	(Zero.) Move the cursor to the left margin of the current line.
[*n*]|	(Vertical bar.) Move the cursor to the beginning of the first column or to the column specified by *n.*
[*n*]w	Move the cursor to the beginning of the next word (or the *n*th word).
[*n*]W	Move the cursor to the beginning of the next word (or the *n*th word), ignoring punctuation.
[*n*]b	Move the cursor to the beginning of the preceding word (or the *n*th word).
[*n*]B	Move the cursor to the beginning of the preceding word (or the *n*th word), ignoring punctuation.
[*n*]e	Move the cursor to the end of the current word (or the *n*th word).
[*n*]E	Move the cursor to the end of the current word (or the *n*th word), ignoring punctuation.
f*x*	Move the cursor forward to the next instance of *x,* where *x* is a character.
F*x*	Move the cursor backward to the preceding instance of *x,* where *x* is a character.

Table 3-1 Motion commands in `vi` (*continued*)

Command	Description
t*x*	Move the cursor forward to one character position before the next instance of *x*, where *x* is a character.
T*x*	Move the cursor backward to one character position after the preceding instance of *x*, where *x* is a character.
[*n*]G	Move the cursor to the specified line number (as in *Go to line number*). Use G alone to move the cursor to the end of the file; use 1G to move the cursor to the beginning of the file.
)	Move the cursor to the beginning of the next sentence (defined as ., !, or ? followed by two spaces or a newline character).
(	Move the cursor to the beginning of the current sentence.
}	Move the cursor to the beginning of the next paragraph. See "Moving by Text Block" for information on how a paragraph is defined.
{	Move the cursor backward to the beginning of a paragraph. See "Moving by Text Block" for information on how a paragraph is defined.
]]	(Right bracket, typed twice.) Move the cursor to the beginning of a new section. See "Moving by Text Block" for information on how a section is defined.
[[	(Left bracket, typed twice.) Move the cursor backward to the beginning of a section. See "Moving by Text Block" for information on how a section is defined.
%	Move the cursor to the matching parenthesis or brace. If you type % when the cursor is not on a parenthesis or brace, `vi` searches forward until it finds one on the current line and then jumps to the matching one.
[*n*]H	Move the cursor to the top-left position on the screen (Home) or the *n*th line from the top of the screen.
[*n*]L	Move the cursor to the bottom-left of the screen (Last) or the *n*th line from the bottom of the screen.
M	Move the cursor to the beginning of the middle line on the screen (Middle).
``	(Back quote key typed twice.) Move the cursor back to where it was before the last absolute motion command. Absolute motion commands are those that move to a precise place (such as a line number or the word you searched for), not a place relative to the cursor position (such as CONTROL-D or 12j).

Your file may have tab (CONTROL-I) characters in it. These characters are represented as several spaces expanding to a tab stop, where the default tab stop is eight spaces. When at a tab, the cursor is on the last of the spaces representing that tab. Try moving the cursor back and forth over tabs to understand how this works.

On rare occasions, your file may contain nonprinting characters. These characters appear as a two-character code, with ^ as the first character; these two characters are treated as a single character.

Moving by text block

The parentheses, (and), move the cursor to the beginning of the previous and next sentences, respectively. A sentence ends with a period, question mark, or exclamation mark, followed by either the end of the line or *two* spaces.

The braces, { and }, move the cursor over paragraphs. A paragraph begins after each empty line, at the `troff` request .bp, or at a paragraph macro. By default, this option uses mm's paragraph macros (that is, the .P and .LI macros). You can change this default with the `set paragraphs` command; see Chapter 4, "Using the ex Line Editor." Each paragraph boundary is also a sentence boundary. You can precede the sentence and paragraph commands with a number to operate over groups of sentences and paragraphs.

Use [[and]] (left and right brackets, typed twice) to move forward and backward over sections, respectively. Sections begin after each macro set in the `sections` option (usually .H and .HU) and each line with a form feed (CONTROL-1) in the first column. Section boundaries are always line and paragraph boundaries.

Moving to a specific line in a file

Type G to move the cursor to the end of your file. The vi editor displays a tilde character (~) on each line past the last line of your text.

Type CONTROL-G to find out the current line number. In response, vi prints, on the last line of your screen, a message telling the name of the file you are editing, the current line number, the total number of lines in the buffer, and a percentage indicating how far

through the buffer you are (in lines). You can also find your current line number by typing the colon character (:) to move to the ex command line and typing

. =

To return to the previous position in the file, type a back quote twice ("). See "Selecting Lines Within a File" in Chapter 4, "Using the ex Line Editor."

You can also move to a specific word or phrase in a file. To search for a particular string in your file, type the slash character (/) to move to the ex command line. You will see the slash character on the bottom line of your screen, waiting for you to specify the string you want to find.

Type

/ *string*

For more information about searching, see "Regular Expressions and Searching" in Chapter 4, "Using the ex Line Editor."

Marking text

To set up a special address use the m*x* command (where *x* is a letter between a and z) and then use this address anywhere you would use a vi address. After you have marked a line, you can refer to it by typing

´*x*

where *x* is the name you gave it.

Scrolling and paging through a file

The vi editor has several commands that allow you to scroll up or down through the text of a file. You cannot combine these commands with other commands such as d (delete) or c (change). You can precede them with a number, *n*, indicating how many lines you want to move. The syntax for this is

[*n*] *scroll-command*

for example,

3 CONTROL-B

You run these commands by holding down the key labeled CONTROL while pressing the indicated letter. Table 3-2 illustrates each keyboard command and its resulting action.

Table 3-2 Scrolling commands in vi

Command	Description
[n] CONTROL-D	Move the cursor down half a screen (or by n lines).
[n] CONTROL-U	Move the cursor up half a screen (or by n lines).
[n] CONTROL-F	Move the cursor to the next full screen (or n full screens forward).
[n] CONTROL-B	Move the cursor to the previous full screen (or n full screens backward).
[n] CONTROL-E	Display another line at the bottom of the screen (or n lines).
[n] CONTROL-Y	Display another line at the top of the screen (or n lines).
[n] CONTROL-P	Move the cursor to the previous line (or n lines backward, same column).
[n] RETURN	Move the cursor to the next line (or n lines forward, same column).
[n] z RETURN	Display the current full screen (or the full screen starting with n).

Inserting text

The commands listed in Table 3-3 are used to insert text. Note that you must press ESCAPE to stop inserting text. Pressing ESCAPE returns you to command mode. You can always undo your last change by typing u in command mode.

◆ **Note** Inserted text may contain newline characters. ◆

Table 3-3 Insert commands in `vi`

Command	Description
a [*text*] ESCAPE	Insert text immediately after the cursor (append).
A [*text*] ESCAPE	Insert text at the end of the current line.
i [*text*] ESCAPE	Insert text immediately before the cursor (insert).
I [*text*] ESCAPE	Insert text at the beginning of the current line.
o [*text*] ESCAPE	Open a new line after the current line and insert text there (open).
O [*text*] ESCAPE	Open a new line before the current line and insert text there.

The `i` command places text to the left of the cursor; the `a` command places text to the right. The `I` command inserts text at the beginning of the line, and the `A` command inserts text at the end of the line. Insert and append a few times to make sure you understand how these commands work. Press ESCAPE to make sure you are in insert mode.

You will often want to add new lines. Type `o`, which creates (opens) a new line after the line containing the cursor. The `O` command creates a new line before the line containing the cursor. After you create a new line, everything you type is inserted on the new line. Press ESCAPE to stop inserting.

To type more than one line of text, press RETURN to end a line. This creates a new line and you continue typing. You can also set the `wrapmargin` option, which automatically moves your cursor to the following line once you have moved to a certain column. See Chapter 4, "Using the `ex` Line Editor," for more information.

Correcting text as you insert

The following characters correct text as you insert it (that is, while you're still in insert mode):

erase   The system erase character (often the ASCII backspace sequence CONTROL-H, DELETE, or #). Deletes the last input character.

end	The system end character (often CONTROL-U, CONTROL-X, or @). Deletes the current input line.
CONTROL-W	Deletes the last word entered.

Your system delete character erases all the input on the current line.

Pressing CONTROL-H or your own erase character erases the last character you typed. Pressing CONTROL-W erases a whole word and leaves the cursor after the preceding word.

Deleting text

The commands listed in Table 3-4 are used to delete text. You can precede most of these commands with a number indicating the extent of the command. For example,

```
3x
```

deletes three characters. Undo the last change by typing u, and repeat the last command by typing . (dot).

Table 3-4 Delete commands in vi

Command	Description
[*n*] x	Delete the character (or *n* characters), starting at the cursor.
[*n*] X	Delete the character (or *n* characters), backward from the character before the cursor.
D	Delete from the cursor to the end of the line.
[*n*] d*motion*	Delete one (or *n*) occurrences of the specified motion. You can use any of the true motion commands here. (See "Using Motion Commands to Move Within a File," earlier in this chapter, for more information.) For example, 3 dw deletes three words.
[*n*] dd	(d typed twice.) Delete current line (or *n* lines including the current line).

You can transpose characters by x'ing the first character that is transposed and then typing p. For example, to correct the word charatcers, move the cursor to the transposed t, type x and then p. This deletes the t and then puts it in the proper place: characters.

You can also delete text by specifying the line numbers you want to delete; see "Deleting Text" in Chapter 4, "Using the `ex` Line Editor."

See "Recovering Lost Text," later in this chapter, if you have deleted some text and want to get it back.

Changing text

The following commands replace text by simultaneously deleting the existing text and inserting new text. You can also precede these commands with a number, *n,* to indicate the extent of the command. For example, `4r` lets you replace four characters. Type `u` to undo these commands. Table 3-5 summarizes `vi`'s replace commands.

Table 3-5 Replace commands in `vi`

Command	Description
`rx`	Replace the character at the cursor with *x.* This is a one-character replacement. You don't need to press the ESCAPE key to end the command.
R [*text*] ESCAPE	Overwrite the characters on the screen with *text.* After you type R, whatever you type overwrites the existing text until you press ESCAPE.
[*n*] s [*text*] ESCAPE	Substitute character (or *n* characters) beginning at the cursor. $ appears at the *n*th position in the text, so you know how much you are changing. Press the ESCAPE key to enter the proper mode.
[*n*] S [*text*] ESCAPE	Substitute the entire current line (or *n* lines). $ appears at end of the current line, or *n* lines are deleted before insertion begins. Press the ESCAPE key to enter the proper mode.
[*n*] c *motion* [*text*] ESCAPE	Change *motion* to *text* where *motion* is a motion command, for example, w for word(s), } for paragraph(s),) for sentence(s), and so on. You can also precede the commands with a number, *n,* to indicate the extent of the command. For example, 4 cw*text* lets you change four words and replace them with *text.* Press the ESCAPE key to enter the proper mode.
[*n*] cc [*text*] ESCAPE	(c typed twice.) Change entire line (or *n* lines). Press ESCAPE to enter the proper mode.
C [*text*] ESCAPE	Change from the cursor to the end of the line.

The `vi` editor prints a message on the last line of the screen telling you how many lines you changed. It also tells you when a change affects text you cannot see.

You can also change text by specifying line numbers on the `ex` command line; see "Changing Text" in Chapter 4, "Using the `ex` Line Editor."

Combining operators and motions

You make larger changes by combining operators (`d` for delete, `c` for change, `s` for substitute, and so on) with the motion commands introduced earlier: `w` for word(s), `)` for sentence(s), `}` for paragraph(s), `/pattern` for context search(es), and so on. The syntax for the general case is

operator motion-command

Move to the beginning of a word and type `dw` to delete a word. Now try `db`; this deletes the word *to the left* of the cursor. The command

`d}`

deletes the text from your current cursor position to the next paragraph delimiter, blank line, or `nroff/troff` command for list or paragraph.

The command

`d)`

deletes the rest of the current sentence. Similarly, `d(` deletes the line to the left of the cursor. The `d(` command deletes the preceding sentence if you are at the beginning of the current sentence, or the current sentence up to the cursor position if you are not at the beginning of the sentence.

You can also use these operators with the `/` (or `?`) search command to change similar phrases in a document. For more information, see "Regular Expressions and Searching," later in this chapter.

Another useful operator is `c` (for change). Use `cw` to change a single word to the text you insert. Press ESCAPE to end. Move to the beginning of a word and type

`cwnew-word`

(followed by ESCAPE). Notice that the end of the text to change was marked with the dollar sign character (`$`).

The `f` and `t` commands are useful with operators such as `c` (change) and `d` (delete) when you want to change a section of text that is not recognized as a delimited word.

Undoing the last command

If you make an incorrect change (whether large or small), use the `u` (undo) command to undo it. Notice that `u` also undoes the previous `u`. Table 3-6 summarizes `vi`'s undo commands.

Table 3-6 Undo commands in `vi`

Command	Description
u	Undo the last command, including a preceding undo command.
U	Undo changes to the current line.

The undo command reverses only a single change. After you make several changes to a line, you may decide that you would rather have the original line back. You can use `U` to restore the current line to its state before you started changing it, but only if you have not moved the cursor to another line before pressing `U`.

If you have made several changes deleting text, you can use `u` to undo only your last change. You can still recover deleted text, however, even if it is too late to use the `u` command. See "Recovering Lost Text," later in this chapter.

Repeating the last command

Use these commands to repeat the last command given. Table 3-7 summarizes `vi`'s repeat commands.

Table 3-7 Repeat commands in `vi`

Command	Description
`.`	Repeat the last command that changed the buffer.
`n`	Repeat the last / or ? search command (next).
`N`	Repeat the last / or ? search command in the opposite direction.
`[n];`	Repeat the last `f`, `F`, `t`, or `T` command (once or *n* times).
`[n],`	Repeat the last `f`, `F`, `t`, or `T` command in the opposite direction (once or *n* times).
`&`	Repeat the last single substitution.

For more information, see "Changing Text" in Chapter 4, "Using the `ex` Line Editor."

Storing text in named buffers

The editor has a set of buffers named `a` through `z`. If you precede any delete or replacement command with

`"`*a*

(where the double quote character indicates "buffer name" and *a* is any single lowercase character), that named buffer will contain the text deleted by the command. For example,

`"a3dd`

deletes three lines, starting at the current line, and puts them in register `a`.

To put them back, move the cursor to where you want the lines and type

`"ap`

or

`"aP`

that is, "put contents of register `a`."

Copying and moving text

The following commands "yank" text (duplicate it in a buffer) and "put" it at another location in the text. These commands are equivalent to what Macintosh users think of as Copy, Cut, and Paste commands.

◆ **Note** You can use the mouse to copy text but not to paste it. To copy text, highlight it with the mouse and press COMMAND-C; the text is copied to the Clipboard. To paste the text, select the new location with the cursor by using the arrow keys or their equivalents, then press COMMAND-V. ◆

Table 3-8 summarizes `vi`'s yank and put commands. In this table, *buf-spec* is the `"a` buffer notation, as specified in "Storing Text in Named Buffers," earlier in this chapter.

Table 3-8 Yank and put commands in `vi`

Command	Description
[*n*] [*buf-spec*] y *motion*	Yank the specified object (word, paragraph, and so on) or *n* objects into a buffer.
[*n*] [*buf-spec*] yy	Yank the current line (or *n* lines) into a buffer.
[*n*] [*buf-spec*] Y	Equivalent to yy.
[*buf-spec*] p	Put the contents of the buffer in the text after the cursor. Lines you yank are placed on new lines following the current line. Other objects, such as words or paragraphs, are inserted immediately following the cursor.
[*buf-spec*] P	Put the contents of the buffer in the text before the cursor. Lines you yank are placed on new lines preceding the current line. Other objects, such as words or paragraphs, are inserted immediately preceding the cursor.

The `vi` editor has a single unnamed buffer in which it saves the last text you deleted or changed, and a set of named buffers (a through z) where you can save copies of text so that you can move text in your file and between files. For more information on these named buffers, see "Storing Text in Named Buffers," earlier in this chapter.

Use the `y` command to place a copy of the specified text into the unnamed buffer. For example, `y3w` puts three words in the buffer.

You can then use the `p` and `P` commands to put the text back in the file. The `p` command puts the text after or below the cursor, while the `P` command puts the text before or above the cursor.

If the text you yank is part of a line or partially spans more than one line, the text is put back after the cursor (or before it, if you use `P`). If the yanked text is whole lines, they are put back as whole lines, without changing the current line. This acts much like an `o` or `O` command.

Try `YP`. This makes a copy of the current line and places it before the current line. The `Y` command is a convenient abbreviation for `yy`. Use `Yp` to copy the current line and place it after the current line. You can give `Y` a count of lines to yank and duplicate several lines. Try typing `3YP`.

You can also copy and move text by specifying line numbers on the `ex` command line. See "Copying and Moving Text" in Chapter 4, "Using the `ex` Line Editor."

Recovering lost text

In addition to the named buffers (`a-z`) and the unnamed buffer (the "undo" buffer), there are nine numbered buffers where the editor places each piece of text you delete (or yank).

The most recent deletion (or yank) is in the undo buffer and also in buffer 1. The next most recent deletion or yank is in buffer 2, and so on. Each new deletion pushes down all older deletions; those older than 9 disappear. If you delete lines and then regret it, you can get the nth previous deleted text back in your file using the command

`"` n`p`

where n is register 1 through 9. The double quote character (`"`) here means "buffer number," n is the number of the buffer (use the number 1 for now), and `p` is the put command, which puts text in the buffer after the cursor. If this doesn't bring back the

text you wanted, type `u` to undo this, and then a period (`.`) to repeat the put command. In general, the period (`.`) repeats the last change you made. When the last command refers to a numbered text buffer, `.` increments the number of the buffer before repeating the command. For example, typing

```
"1pu.u.u.u.u.u.u.
```

shows you all the deleted text that has been saved (nine deletions). You can omit the `u` commands here to gather all this text in the buffer, or you can stop after any command to keep only the text recovered so far. You can use `P` instead of `p` to put the recovered text before rather than after the cursor.

You can use the `ex` commands `co` and `m` to copy and move text, respectively. For more information, see Chapter 4, "Using the `ex` Line Editor."

Regular expressions and searching

You can search for words or phrases in files by typing a slash (`/`), followed by the word or phrase you want to find, followed by RETURN. For example, type

```
/empire RETURN
```

to find the next instance of `empire` in the file. This is an `ex` command. When you type the slash, the cursor moves to the `ex` command line.

Regular expressions use special characters and notation to specify a set of character strings. For example, the regular expression `.` (dot) matches *any* single character. Therefore, the search command

```
/c.t
```

takes you to the next occurrence of all words containing a `c` followed by any character followed by a `t` (`cat`, `cbt`, `cct`, and so on).

Regular expressions can be extremely useful when you're editing a file. For more information about them, see "Regular Expressions and Searching" in Chapter 4, "Using the `ex` Line Editor."

Working with multiple files

You can edit more than one file at the same time. To open more than one file at once, enter `vi` with the command

`vi` *file1 file2*

file1 is opened first. When you write the file, you can move to the second file by typing the colon character (`:`) to move to the `ex` command line and typing

`n`

You must write the first file before trying to get into the second file, or `vi` won't let you move.

To return to the first file, press

CONTROL-`^`

(On many terminals you must press CONTROL-SHIFT-6 to get this sequence.)

You can also read the contents of another file into the buffer by typing the colon character (`:`) to move to the `ex` command line and typing

`r` *filename*

A copy of *filename* is inserted after the current line. See "Copying Another File to the Current Buffer" in Chapter 4, "Using the `ex` Line Editor," for more information.

For more information about editing several files at the same time, see "Working With Multiple Files" in Chapter 4, "Using the `ex` Line Editor."

Using shell commands in `vi`

There are several ways of interacting with the shell from within `vi`. All these require that you type the colon character (`:`) to move to the `ex` command line and then type the appropriate `ex` command. See "Using Shell Commands in `ex`" in Chapter 4, "Using the `ex` Line Editor."

Setting options

You can customize your `vi` environment by setting various `ex` options. To set an option in `vi` type the colon character (`:`) to move to the `ex` command line and then type the appropriate `set` command. Table 3-9 summarizes some `ex` command options that are particularly useful in `vi`.

Table 3-9 Summary of `ex` command options

Option	Description
set wi[ndow]=*n*	Change the number of lines in your editing window.
set scr[oll]=*n*	Change the number of lines you scroll through with the CONTROL-D command.
set para[graphs]=*macro-strings*	Change the strings `vi` searches for when you press `{` or `}`. Valid strings are the paragraph and list macros from the `mm` and `ms` macro packages.
set sect[ions]=*macro-strings*	Change the strings `vi` searches for when you press `[[` or `]]`. Valid strings are the section macros from the `mm` and `ms` macro packages.
set redraw	Force a dumb terminal to redraw the characters to the right of the cursor as you enter text in `vi`.
set w[rap]m[argin]=*n*	Set the column where the cursor automatically returns to the left margin.
set nomesg	Prevent messages from other people while in `vi`.

For more information about setting options and a list of all the available options, see "Setting Options" in Chapter 4, "Using the `ex` Line Editor."

Mapping and abbreviations

Mapping and abbreviations are available as a joint facility of `ex` and `vi`. The mapping or abbreviation must be defined on the `ex` command line, but it is useful only in visual mode. Since `vi` is required for these to work, the examples are structured to work within `vi`. These commands require that you type the colon character (`:`) to move to the `ex` command line before typing the appropriate command.

Preventing nonprinting characters from being interpreted as commands

If you need to enter a control character, such as the system *erase* or *end* character, as part of your file, you need to precede it with an escape character so that `vi` will not interpret it as a command. You can do this as follows:

\	If the character or sequence prints on your screen, for example `$`, but has a special meaning you want to override, precede the character with a backslash.
CONTROL-V	If the character or sequence is invisible on your screen (for example, CONTROL-1), precede it with CONTROL-V.

For example, to type an escape character in your file, type CONTROL-V before you press the ESCAPE key. The CONTROL-V prints a caret (`^`) at the current cursor position, indicating that the editor expects you to type a control character. When you press the ESCAPE key, you see `^[` in your file. After pressing CONTROL-V, you can type any nonprinting character except the null (`@`) character or the line-feed (CONTROL-J) character. Using CONTROL-V is the only way to insert CONTROL-S or CONTROL-Q.

Using the `map` command to create macros

In `vi` you can use the `ex` command `map` to create a macro. A macro sets a string (usually a single key) equal to a command or sequence of commands.

The general format of the `map` command in `vi` is

`:map` *string definition*

(followed by RETURN), where *string* is usually a single keystroke or function key, and *definition* is the map definition representing a command or sequence of commands.

After you enter a `map` command, type *string* to perform the specified command; *string* can't be more than 10 characters and *definition* can't be more than 100 characters. If it takes longer than a second to type *string*, however, `vi` times you out before recognizing the string (you can prevent this by setting the option `notimeout`; see "Setting Options" in Chapter 4, "Using the `ex` Line Editor").

When you set up a `map` from within `vi`, it lasts only as long as your current editing session. If you're editing several files in one session (by entering several names on the command line) and you use the `:n` command to edit them, however, the `map` definitions hold for all files until you quit the editor. To make a more permanent definition, insert it in the `EXINIT` variable in your `.profile` (or `.login`) file or in `$HOME/.exrc` (see "Setting Options," in Chapter 4, "Using the `ex` Line Editor").

CONTROL-V allows you to insert nonprinting characters (RETURN, blanks, tabs, control sequences) in your `map` definition. Pressing RETURN ends the `map` command, so to include RETURN, ESCAPE, or any other nonprinting character in *definition*, you must precede it with a CONTROL-V escape character.

Suppose you want to create a macro to assign a single keystroke to perform a compound write and quit (wq) command. To map the `q` key, for example, to perform this compound command, type

`:map q :wq` CONTROL-V RETURN

This maps the sequence

`:wq` RETURN

to the character `q`. When you type `q` in `vi` command mode, the command `:wq` RETURN runs. Without CONTROL-V, the first RETURN ends the `map` command, rather than becoming part of the `map` definition. If you use CONTROL-V, however, the first RETURN becomes part of *definition*, and the second ends the `map` command.

You can use `#` in the map command to represent function keys. Some terminals don't have function keys, but most terminals do. If *string* is `#0` through `#9`, you map to the corresponding function key, *not* the two-character sequence `#n`. (Note that on the Apple Extended Keyboard, `#n` refers to the corresponding key on the numeric key pad.)

You can also use the `map` command so that one key calls a second key, which calls the first one again. This is useful to repeat an editing action throughout an entire file. When you're writing these double maps that call each other, however, make sure the pattern you are searching for changes as each command runs.

When you `map` a key to the function, choose a key that is not already associated with a function you need to use. For example, suppose you have a file containing names and telephone numbers in one format, and you want to change it globally as follows:

Initial format   *New format*

```
. . .                           . . .
Lemonwedge, Ned                 Lemonwedge, Ned 555-1234
555-1234                        Spears, Ray 555-4567
Spears, Ray
555-4567

. . .                           . . .
```

To do this, map the keys `g` and `h` as follows:

:map g /^[0-9]/CONTROL-V RETURN kJh

:map h /^[0-9]/CONTROL-V RETURN kJg

Then, in `vi`'s command mode, typing `g` runs the command

/^[0-9]/CONTROL-V RETURN kJh

That is, the search command looks for a line beginning with a number 0 through 9. You must press RETURN after this command to run it, and you must precede the RETURN in this definition with a CONTROL-V escape character. When the cursor is positioned on a line beginning with a number 0 through 9, the command

kJh

indicates "go up one line" (`k`), "join this line and the next line" (`J`), and "call `h`," which repeats this sequence of commands and calls `g`.

These commands finish when there are no lines left in the file that begin with a number 0 through 9. The number of lines beginning with a number is reduced by one each time the `J` command runs. Otherwise, these commands can go into an infinite loop, calling each other indefinitely. If this happens, send an interrupt.

You can unmap keys with

: unmap *string*

The undo command reverses the function of the entire macro.

Abbreviations

You can define a short string that expands to a longer string in the text. The syntax of the commands to perform abbreviate and unabbreviate (ab and una) is

: *abbreviation-command wd* [*word*]

(followed by RETURN), where *abbreviation-command* is ab[breviate] or una[bbreviate], *wd* is the abbreviation you're defining, and *word* (only applicable for abbreviate) is the string it represents. Note that the order of the arguments is the reverse of what you might expect.

The *wd* argument can't be more than 10 characters long and the *word* argument can't be more than 100 characters long. If it takes longer than a second to type *wd*, however, vi times you out before recognizing the string. (You can prevent this by setting the ex option notimeout; see "Setting Options" in Chapter 4, "Using the ex Line Editor.")

After you enter an abbreviate command, and type wd, it is immediately translated into *word*. For example, suppose you type this command

:ab dtp Developer Technical Publications

Now suppose you enter (in insert mode) the string dtp (followed by spaces, a newline character, or punctuation marks). This string is immediately expanded to the phrase "Developer Technical Publications," which is entered as part of your vi buffer. If you type dtp as part of a larger word, however, it is left alone. The abbreviation echoes as you type it, and when you type a delimiter that sets it apart as a single word, it immediately expands into the longer string.

Avoid using abbreviations that are words in themselves, for example *a* or *is*.

Additional features

The following sections describe some additional features of the `vi` screen editor.

Saving files and quitting `vi`

You should periodically save the text you have entered in the buffer so that you don't accidentally lose any changes you have made. To save the file, use the `ex` command

`:w`

This command saves (or writes) the text into a permanent file, overwriting any previous version of the file. See "Saving Text and Quitting `ex`" in Chapter 4, "Using the `ex` Line Editor."

You can quit `vi` by making sure you are in command mode (press the ESCAPE key) and giving the command

`ZZ`

This command writes the changes you made to the buffer back into the file you were editing and quits `vi`. However, if you did not change to the buffer, this command does not force a write operation.

If you want to force a write operation (this is especially recommended when recovering a file using `vi -r`), use the `ex` command

`:wq`

There are several `ex` commands for exiting `vi` in different ways. To use them type the colon character (`:`) to move to the `ex` command line and then type the appropriate `ex` command. For more information see "Saving Text and Quitting `ex`" in Chapter 4, "Using the `ex` Line Editor."

Troubleshooting

The following sections describe solutions to some common problems you may experience during a `vi` session.

Redrawing the screen

Occasionally, you may need to redraw (or refresh) the screen. (For example, a program can write to the screen, or line noise can jumble up the screen.)

Press CONTROL-L to redraw the screen. On certain terminals, CONTROL-R works instead.

You can redraw the screen so that a specified line shows at the top, middle, or bottom. To do this, move the cursor to that line and type

z

Pressing RETURN after the `z` command redraws the screen with the line at the top; a period (.) after the `z` places the line at the center; a minus sign (-) after the `z` places it at the bottom.

Speeding up a slow system

You can reduce the overhead of refreshing the screen for each change, scroll, and so on by limiting the window size. Use this command to start `vi` with a specific window size

vi -w*n filename*

where *n* is a number less than 23. For example,

vi -w3 *filename*

This sets the initial size of the window to three lines, and allows the window to expand as you add lines.

You can control the size of the window that `vi` redraws when the screen clears by specifying a window size as an argument to any of the following commands:

: / ? [[ ]] ` ´

If you search for a string in a file, preceding the first search command with a small number (for instance, `3`) draws three-line windows around each instance of the string.

You can make the window larger or smaller by giving a number after the `z` command and before the following RETURN, period (`.`), or minus (`-`). For example, the command

```
z5.
```

redraws the screen with the current line in the center of a five-line window. The command `5z.` has an entirely different effect, placing line 5 in the center of a new window.

If `vi` is updating large portions of the display, you can interrupt it by sending an interrupt signal (usually DELETE or CONTROL-C). This may partially confuse `vi` about what is displayed on the screen. You can clear up the confusion by typing CONTROL-l or by moving or searching again, ignoring the current state of the display.

Creating temporary file space

The `vi` editor prints the message

```
Out of temp filespace
```

when it doesn't have enough buffer space (in `/tmp`) to hold the file. It might refuse to open the file when it prints this, or it might open it and load only part of the file into the buffer. The latter is dangerous because if you write the file when only half of it is in the buffer, you can lose the other half.

The best thing to do in response to that message is to quit `vi`. Type the colon character (`:`) to move to the `ex` command line and then type

```
q!
```

Then go up to `/tmp` and delete any files that aren't necessary. After you've deleted some files, go back and try opening your document again.

If your file is very large, you may have to use the A/UX `split` command to break it into smaller text files before you can use `vi` to edit it. See `split(1)` in *A/UX Command Reference* for more information.

Recovering lost files

If the system crashes or you receive a hang-up signal on a dial-in line, you can recover your work even if you did not write the changes to the file. Move to the directory you were in when the system crashed, and give the command

vi -r *filename*

(the -r flag option for recover) where *filename* is the file you were editing when the system went down. This command tells you if there are any files to be recovered.

In some cases, a few lines of the file may be lost. These lines are almost always the last few you changed.

♦ **Note** It is not advisable to exit a file that you have just recovered using the ZZ command, since ZZ does not guarantee that the file will be written. After you have checked the recovered file, you should save the file using the :w command. ♦

Command summary

Tables 3-10 and 3-11 give brief descriptions of editing and inserting commands in vi.

Table 3-10 Summary of `vi` editing commands

Command	Description
[*n*] CONTROL-B	Move the cursor to the previous screen (or *n* screens backward).
[*n*] CONTROL-D	Move the cursor down half a screen (or by *n* lines).
[*n*] CONTROL-E	Display another line at the bottom of the screen (or *n* lines).
[*n*] CONTROL-F	Move the cursor to the next screen (or *n* screens forward).
CONTROL-G	Give current line number and filename and the percentage along in the file.
CONTROL-H	Move the cursor back one space.
[*n*] CONTROL-J	Move down one line in the same column (or *n* lines).
CONTROL-L	Redraw screen (certain terminals only—others use CONTROL-R).
[*n*] CONTROL-M	Same as pressing the RETURN key.
[*n*] CONTROL-P	Move the cursor to the previous line (or *n* lines backward, same column).
CONTROL-R	Redraw screen (most terminals—others use CONTROL-1).
[*n*] CONTROL-U	Move the cursor up half a screen (or by *n* lines).
[*n*] CONTROL-Y	Display another line at the top of the screen (or *n* lines).
CONTROL-^	Edit the alternate file. See "Working With Multiple Files," earlier in this chapter, for more information.
a [*text*] ESCAPE	Insert *text* immediately after the cursor (append).
A [*text*] ESCAPE	Insert *text* at the end of the current line.
[*n*] b	Move the cursor to the beginning of the preceding word (or *n*th word).
[*n*] B	Move the cursor to the beginning of the preceding word (or *n*th word), ignoring punctuation.
[*n*] c *motion* [*text*] ESCAPE	Change *motion* to *text*, where *motion* is a motion command.
cc [*text*] ESCAPE	(c typed twice.) Change entire line (or *n* lines).
C [*text*] ESCAPE	Change from the cursor to the end of the line.
[*n*] d *motion*	Delete one (or *n*) occurrences of the specified motion. You can use any of the true motion commands here. (See "Using Motion Commands to Move Within a File," earlier in this chapter, for more information.) For example, 3dw deletes three words.
[*n*] dd	(d typed twice.) Delete current line (or *n* lines including current line).

Table 3-10 Summary of `vi` editing commands (*continued*)

Command	Description
D	Delete from the cursor to the end of the line.
[*n*] e	Move the cursor to the end of the current word (or *n*th word).
[*n*] E	Move the cursor to the end of the current word (or *n*th word), ignoring punctuation.
[*n*] f*x*	Move the cursor forward to the first instance of *x* (*n* specifies the *n*th instance).
[*n*] F*x*	Move the cursor backward to the first instance found of *x* (*n* specifies the *n*th instance).
[*n*] G	Move the cursor to the specified line number (Go to line number). G alone moves the cursor to the end of the file. 1G moves to the beginning of the file.
h or ←	Move the cursor left a space. (BACKSPACE also works.)
[*n*] H	Move the cursor to the top-left position on the screen (Home) or the *n*th line from the top of the screen.
i [*text*] ESCAPE	Insert *text* immediately before the cursor (insert).
I [*text*] ESCAPE	Insert *text* at the beginning of the current line.
[*n*] j or [*n*] ↓	Move the cursor down a line (in the same column) or down *n* lines.
[*n*] J	Join current line with next line or the next *n* lines.
[*n*] k or [*n*] ↑	Move the cursor up a line (in the same column) or up *n* lines.
[*n*] l or [n] →	Move the cursor right a space or *n* spaces. (Space bar also works).
[*n*] L	Move the cursor to the bottom-left of the screen (Last) or the *n*th line from bottom of the screen.
m	Mark the current position of the cursor in the register specified by the following letter (a through z). Return to this position with ` and the register letter.
M	Move the cursor to the beginning of the middle line on the screen (Middle).
n	Repeat the last / or ? search command (next).
N	Repeat the last / or ? search command in the opposite direction.

Table 3-10 Summary of `vi` editing commands (*continued*)

Command	Description
o [*text*] ESCAPE	Open a new line after the current line and insert *text* there (open).
O [*text*] ESCAPE	Open a new line before the current line and insert *text* there.
[*buf-spec*] p	Put the contents of the buffer in the text after the cursor. Lines you yank are placed on new lines following the current line. Other objects, such as words or paragraphs, are inserted immediately following the cursor. The *buf-spec* argument specifies a buffer "a through "z.
[*buf-spec*] P	Put the contents of the buffer in the text before the cursor. The *buf-spec* argument specifies a buffer "a through "z.
Q	Quit visual mode and go to the `ex` command line.
[*n*] RETURN	Move the cursor to the next line (or *n* lines forward, first column).
[*n*] r*x*	Replace the character (or *n* characters) at the cursor with *x*. This is a one-character replacement. You don't need to press ESCAPE to cancel the command.
R [*text*] ESCAPE	Overwrite the characters on the screen with *text*. After you type R, whatever you type overwrites the existing text until you press ESCAPE.
[*n*] s [*text*] ESCAPE	Substitute the character (or *n* characters) beginning at the cursor. The $ character appears at the *n*th position in the text, so you know how many you are changing.
[*n*] S [*text*] ESCAPE	Substitute the entire current line (or *n* lines). The $ character appears at the end of the current line, or *n* lines are deleted before insertion begins.
[*n*] t*x*	Move the cursor forward to just before the first instance of *x* (or the *n*th instance).
[*n*] T*x*	Move the cursor backward to just before the first instance of *x* (or the *n*th instance).
u	Undo the last command, including a preceding undo command.
U	Undo changes to the current line.
[*n*] w	Move the cursor to the beginning of the next word (or *n*th word).
[*n*] W	Move the cursor to the beginning of the next word (or *n*th word), ignoring punctuation.
[*n*] x	Delete the character (or *n* characters), starting at the cursor.
[*n*] X	Delete the character (or *n* characters), backward from the character before the cursor.

Table 3-10 Summary of `vi` editing commands (*continued*)

Command	Description
[*n*] [*buf-spec*] y *motion*	Yank the specified object (word, paragraph, and so on) or *n* objects into a buffer. The *buf-spec* argument specifies a buffer "a through "z.
[*n*] [*buf-spec*] yy	Yank the current line (or *n* lines) into a buffer. The *buf-spec* argument specifies a buffer "a through "z.
[*n*] [*buf-spec*] Y	Equivalent to yy. The *buf- spec* argument specifies a buffer "a through "z.
[*n*] z	Display the current full screen (or the full screen starting with *n*). Change the placement of the current line by following z with one of these characters:
	RETURN Place current line at the top of the screen.
	. Place current line at the center of the screen.
	– Place current line at the bottom of the screen
ZZ	Quit vi, performing a write operation first if changes were made to the file.
–	Move the cursor to the beginning of the preceding line. Scroll if necessary.
+	Move the cursor to the beginning of the next line. Scroll if necessary.
[*n*] $	Move the cursor to the end of the current line, or to the end of the line *n* lines forward in the file.
^	Move the cursor to the beginning of the first word on the line.
0	(Zero.) Move the cursor to the left margin of the current line.
[*n*] |	(Pipe or vertical bar.) Move the cursor to the beginning of the first column or to the column specified by *n*.
)	Move the cursor to the beginning of the next sentence (defined as ., !, or ?) followed by two spaces or a newline character.
(	Move the cursor to the beginning of the current sentence.
}	Move the cursor to the beginning of the next paragraph (the default definition of a new paragraph is .P, .LI, or .bp) or to the next blank line.
{	Move the cursor backward to the beginning of a paragraph (the default definition of a new paragraph is .P, .LI, or .bp) or to the last blank line.

(continued)➡

Table 3-10 Summary of `vi` editing commands (*continued*)

Command	Description
]]	(Right bracket, typed twice.) Move the cursor to the beginning of a new section (default definition is by `.H` or `.HU`).
[[	(Left bracket, typed twice.) Move the cursor backward to the beginning of a section (default definition is by `.H` or `.HU`).
%	Move the cursor to the matching parenthesis or brace. If you type % when the cursor is not on a parenthesis or brace, `vi` searches forward until it finds one on the current line and then jumps to the matching one.
`"`	(Back quote key typed twice.) Move the cursor back to where it was before you used the last absolute motion command. Absolute motion commands are those that move to a precise place.
.	Repeat the last command that changed the buffer.
[*n*] ;	Repeat the last `f`, `F`, `t`, or `T` command (once or *n* times).
[*n*] ,	Repeat the last `f`, `F`, `t`, or `T` command in the opposite direction (once or *n* times).
&	Repeat the last single substitution.

Table 3-11 summarizes the commands used in `vi`'s insert mode.

Table 3-11 Summary of `vi` insert commands

Command	Description
CONTROL-D	During an insert, tab backward over `autoindent` white space at the beginning of a line.
CONTROL-I	Input a tab.
CONTROL-Q	Precede a single character or control sequence with the escape character CONTROL-Q.
CONTROL-T	Insert a shift-width wide white space if pressed at the beginning of a line with `autoindent` set.
CONTROL-V	Precede a single character or control sequence with the escape character CONTROL-V.
CONTROL-W	Delete the last word entered.
erase	The system erase character (often DELETE, CONTROL-H, or #). Delete the last input character.
delete	The system end character (often CONTROL-U, CONTROL-X, or @). Delete the current input line.

4 Using the ex Line Editor

This chapter provides a detailed description of the commands and capabilities of the ex line editor.

What is `ex`?

The `ex` editor is an advanced version of the `ed` line editor. Although `ex` is equipped with many powerful commands, most UNIX users prefer to use `vi` because `ex` doesn't maintain an updated text display on the screen. Even if you never use `ex` on its own, it is important to learn about `ex` commands and their functions because most `ex` commands are readily accessible from `vi`. For more information, see Chapter 3, "Using the `vi` Screen Editor."

Once you have started `ex`, you enter commands at the `ex` prompt, the colon (`:`). In the `ex` editor, commands are words (such as `write` or `edit`), which you can abbreviate. A complete list of commands and their abbreviations appears at the end of this chapter.

Starting `ex`

There are two ways to start `ex`. The first is to start `ex` from a CommandShell window. To start `ex` in this manner, type

`ex` *filename* `. . .`

where *filename* is the name of the file to edit. This is the fastest way to start `ex`.

You can also use the Commando command line interface to open a new file. To do this, open the `ex` Commando dialog box, click the "Enter name of new file" button, type the name of the file, and then click OK. Click the "ex" button to open the new file.

The `ex` application starts. You return to the CommandShell window, and the empty file you created is opened. See *A/UX Essentials* for more information about Commando.

Syntax and initialization

The following sections describe the command syntax and initialization for the `ex` editor.

Command syntax for the `ex` line editor

The command syntax for `ex` is as follows:

`ex [-] [-v] [-t [tag]] [-r [filename]] [-l] [-wn] [-x] [-R] [+command]`
`[filename ...]`

Table 4-1 summarizes the flag options for the `ex` line editor.

Table 4-1 `ex` flag options

Option	Description
–	(Minus sign.) Suppress interactive feedback. This is useful when you write shell scripts that use the `ex` editor.
-v	Equivalent to using `vi`.
-t [tag]	Start editing the file at *tag* (usually a spot marked with the `ctags` program). Equivalent to an initial `tag` command. See "Editing Programs," later in this chapter.
-r [filename]	Recover a file after an editor or system crash; if you don't specify *file*, `-r` lists the saved files.
-l	Set the `showmatch` and `lisp` options for editing LISP programs. See "Setting Options," later in this chapter.
-wn	Set the window size to *n* lines.
-x	Prompt for a key to encrypt and decrypt the file (see `crypt`(1) in *A/UX Command Reference*). The file should already be encrypted using the same key.
-R	Set the `readonly` option, making it impossible to write the file with the `write` command.
+ [command]	Move to the line specified by *command* where *command* is either a regular expression (see "Regular Expressions and Searching") or a line number (for example, `+100` starts editing at line 100). If you omit *command*, `ex` moves the cursor to the last line of the first file.
filename	The file(s) to edit.

Initialization procedures of the `ex` editor

When you start `ex`, it sets up your editing environment with the following steps:

- reads the `TERM` variable to find out what terminal you're using
- sets any options you've specified in the `EXINIT` environment variable (usually set in the `.profile` (or `.login`) file in your home directory)
- sets any options you've specified in the `.exrc` file in the current directory or your home directory

You can set the same options with either the `EXINIT` environment variable or `.exrc` files. The options are described in "Setting Options," later in this chapter.

Opening a file

To create and open a new file (or open an existing file) in `ex`, type

`ex` *filename*

where *filename* is the name of the file you're creating (or opening). For example, to open the file `london`, type

`ex london`

When you use `ex` to create a new file, `ex` opens some temporary storage space that is called the **buffer.**

When you edit an existing file, `ex` places a copy of that file in the buffer. Changes you make to the text in the buffer (for example, to a copy of `london`) are made only to this temporary copy. The contents of the file are not changed until you save the file. See "Saving Text and Quitting `ex`," later in this chapter.

▲ **Warning** You should periodically save your changes to the file to prevent losing material if the system crashes or is interrupted. ▲

The different modes of `ex`

The `ex` editor has a number of modes:

- **Command mode** is `ex`'s primary mode; that is, when you invoke `ex` you are initially in command mode and everything you type is interpreted as a command. In command mode, you enter commands at the colon (`:`) prompt and end them by pressing RETURN.

- In **input mode,** `ex` assumes that what you type is text (rather than commands), and it doesn't display a prompt. Invoking the `append`, `insert`, or `change` commands in command mode places you in input mode. Resume command mode by typing a period (`.`) at the beginning of a line and pressing RETURN.

- You can also enter open mode from `ex` using the `o` command. **Open mode** allows you to use `vi` commands, but limits your movement to within one line at a time. (It is like visual mode with a screen one line long.) This is convenient if you want to use `vi` commands on a dumb terminal. (A smart, or addressable-cursor, terminal is required for you to use `vi` itself.) Type `Q` to return to `ex`.

- When you invoke `ex` with the command `-v` you enter `vi`. **Visual mode** allows movement and editing throughout the displayed screen of text. See "Switching to `vi`."

Switching to `vi`

When you invoke `ex` you can switch to `vi` by giving the command

`vi`

at the `ex` colon prompt on the bottom line of your screen. This invokes visual mode and places the current line as the first line on the screen. To return to `ex` from visual mode, type

`Q`

Using special keys in ex

The following keys have special meaning in A/UX.

RETURN   The RETURN (carriage return) key cancels all ex commands.

                  The examples in this chapter assume that you press RETURN after all
                  commands unless shown otherwise.

The interrupt   The interrupt key sequence (set to CONTROL-C in A/UX) sends an
key sequence   interrupt signal to the editor. It is a forceful way of canceling a
                  command after you have pressed RETURN.

Displaying text and selecting lines within a file

In ex you can display text on the screen by specifying a line number or range of line
numbers followed by the print command. The next section describes how to specify
which lines you want to select.

Selecting lines within a file

You can prefix most commands with addresses. These addresses tell ex which lines to
perform the command on. For example, 10print prints the tenth line in the buffer.

Here are a few basic rules to follow when using line addresses in ex:

- Commands that don't require an address (such as the quit command to leave
  ex) regard an address as an error. "Command Summary," at the end of this chapter,
  specifies which commands need addresses and which don't.

- For ex commands that require an address, ex assumes a default address if you
  don't supply one. If you give more addresses than a command requires, ex uses
  the last one or two, depending on the command being attempted.

- For `ex` commands requiring two addresses, the second address must follow the first address. If you use two addresses, you can separate them with a comma (`,`) or a semicolon (`;`). Using `,` calculates both addresses relative to the current line. For example, if you are on line 1,

  ```
  +2,+4print
  ```

  prints lines 3 through 5. Using `;` calculates the second address using the first address as the current line. For example, if you are on line 1,

  ```
  +2;+4print
  ```

  prints lines 3 through 7.

Throughout the rest of this chapter, the term *lineno* (line number) denotes an expression that identifies a single line in a file, numbered with *lineno*. This includes a search for a pattern (see `/pattern[/]`, below).

Throughout the rest of this chapter, the term *line-spec* (line specifier) denotes an expression that identifies zero or more lines in a file. A valid *line-spec* can be a single line number *lineno;* a range of line numbers *line1, line2;* a context search resulting in zero or more lines (see "Regular Expressions and Searching"); or a regular expression resulting in zero or more lines (see "Regular Expressions and Searching"). Commonly used abbreviations include `$` (the last line of a file) and `.` (the current line). In the following commands, *line-spec* defaults to the current line unless stated otherwise.

A *line-spec* may consist of any one of the following expressions:

`.`	Dot (`.`) indicates the current line. This is the default address for most commands. Most commands leave the current line as the last line they affect.
line1, line2	A range of line numbers beginning with *line1* and ending with *line2*.
lineno	Move to *lineno*. You can find out the current line number by typing a `.=` sequence.
`$`	The last line in the buffer.
`%`	An abbreviation for `1,$`; the entire buffer.
`+[n]`	Forward *n* lines from the current line. `+3` and `+++` are equivalent; if the current line is line 100, they address line 103.
`-[n]`	Backward *n* lines from the current line. `-3` and `---` are equivalent; if the current line is line 103, they address line 100.

/*pattern*[/]	Forward to a line containing the regular expression *pattern* (see "Regular Expressions and Searching," later in this chapter). ex searches forward until it reaches the end of the file, then it searches from the first line of the file to where you began your search. If you want to print the next line containing *pattern*, you don't have to include the trailing /. / is shorthand for "search forward for the last pattern you scanned for."
?*pattern*[?]	Backward to a line containing the regular expression *pattern* (described under "Regular Expressions and Searching"). ex searches backward until it reaches the beginning of the file, then it searches from the last line of the file to where you began your search. If you want to print the next line containing *pattern*, you don't have to include the trailing ?. ? is shorthand for "search backward for the last regular expression you scanned for."
"	(Back quote typed twice.) Refers to your position before the last absolute motion (an **absolute motion** specifies the line to move to, while a **relative motion** specifies the distance to move from the current line.)
'x	(Where *x* is a letter from a to z.) Refers to a location you marked with the mark command. This is described in "Marking Text," later in this chapter.

You can also add a number to the end of a command to specify the number of lines involved. For example, d5 deletes five lines, starting with the current line. If the number specified is larger than the number of lines between the current line and the end of the file, ex performs the operation to the end of the file.

Using motion commands to move within a file

ex provides a number of ways to move through your file, and these are collectively referred to as motion commands. Table 4-2 summarizes ex's motion commands.

Table 4-2 ex motion commands

Command	Description
lineno	Move to *lineno*.
$	Move to the last line in the buffer.
+ [*n*]	Move forward from the current line. If followed by *n*, it means to move to the start of the line *n* lines forward in the buffer.
− [*n*]	Move backward from the current line. If followed by *n*, it means to move to the start of the line *n* lines backward in the buffer.
/ *pattern* [/]	Move forward to a line containing the regular expression *pattern* (described under "Regular Expressions and Searching").
/	Move forward using the last regular expression scanned for.
/ /	Move forward using the last regular expression used in a substitution (described under "Changing Text").
? *pattern* [?]	Move backward to a line containing the regular expression *pattern* (described under "Regular Expressions and Searching").
?	Move backward using the last regular expression scanned for.
? ?	Move backward using the last regular expression used in a substitution (described under "Changing Text").
"	(Back quote typed twice.) Move to the location before the last absolute motion.
' x	(A single back quote followed by a lowercase letter a through z.) Move to a location you marked with the mark command (described in "Marking Text").
CONTROL-D	Move forward half a screen. You can change this with the set scroll option, described under "Setting Options."
RETURN	Move to the next line.

Determining line appearance

Three commands control how text displays on the screen: print, list, and number.

The print command sets the default printing style. It displays nonprinting characters as ^*x* and delete (octal 177) as ^?. The format of the print command is

[*linespec*] p[rint] [*n*]

It prints the current line, the lines specified by *linespec*, or the next *n* lines. You can also add a `p` to the end of many commands to print the current line after the command completes.

The `list` command displays the specified lines with tabs indicated with `^I` and the ends of lines indicated with a `$`. The format of the `list` command is

[*linespec*] l[ist] [*n*]

It displays the current line, the lines specified by *linespec*, or the next *n* lines. You can also add an `l` to the end of many commands to list the current line after the command completes.

The `number` command prints the line number before each specified line. The format of the `number` command is

[*linespec*] nu[mber] [*n*]

or

[*linespec*] # [*n*]

It prints the current line, the lines specified by *linespec*, or the next *n* lines. You can also add a `#` to the end of many commands to number the current line after the command completes.

See "Setting Options" for more information.

Determining the appearance of the current line on the screen

The `z` command determines where the current line appears on the screen. (If you prefix the command with a line number *lineno*, that line number becomes the current line.) There are several different forms of this command.

[*lineno*] *z*	Print the next full screen of lines with the current line at the top of the screen.
[*lineno*] *z*+	Print the next full screen of lines with the current line at the top of the screen.
[*lineno*] *z*–	Print the screen with the current line at the bottom.

[*lineno*] z.	Print the screen with the current line at the center.
[*lineno*] z=	Print the screen with the current line in the center, surrounded with lines of – characters.
[*lineno*] z^	Print the screen two windows before the current line.

Inserting text

The two basic `ex` commands for adding text to your file are the `append` and `insert` commands.

The `append` command adds text *after* the specified line. The command syntax is

[*lineno*] a[ppend] RETURN [*text* RETURN] . RETURN

where *lineno* is the line number, *text* is the text you enter. To append text to the start of the buffer, use the command `0a` (this appends to line 0). A variant form, the `append!` command, changes the setting of the `autoindent` option while appending (described under "Setting Options").

The `insert` command adds text *before* the specified line. The command format is

[*lineno*] i[nsert] RETURN [*text* RETURN]. RETURN

where *lineno* is the line number, *text* is the text you enter. A variant form, the `insert!` command, changes the setting of the `autoindent` option during an insert (described under "Setting Options").

Deleting text

The `delete` command removes the specified lines. You delete a specified range of line numbers using the format

[*linespec*] d[elete] [*n*]

For example, you can delete lines 5 through 20 of your file with the command

5,20d

You can delete line 5 with the command

```
5d
```

or you can specify the line you want to delete using a regular expression. For example, if the line you want to delete ends in the word *finish,* use the command

```
/finish$/d
```

You can also delete *n* lines with the command

```
d[elete] n
```

For example,

```
d3
```

Changing text

The `change` command replaces existing text with new text. The command format is

[*line1*[, *line2*]]c[hange][*n*] RETURN *text* RETURN

This changes either *line1*, the range *line1* through *line2* inclusive, or, when *n* is specified, the next *n* lines. A variant form, the `change!` command, changes the setting of the `autoindent` option during the change. (This option is described under "Setting Options.") The command format is

[*line1*[, *line2*]]c[hange]![*n*] RETURN *text* RETURN.

You can also change a string with the `substitute` command. The simplest format of this command is

s/*pattern*/*replacement*[/]RETURN

This replaces the first instance of *pattern* with the *replacement* on the current line. Regular expressions are used commonly in substitutions. For example, if you had the following line in your file

```
Mail The Cleaning Bill To My Gettysburg Address
```

and you typed

```
s/Mail/Send/
```

it would change your line to

```
Send The Cleaning Bill To My Gettysburg Address
```

Substitutions are also commonly used when you want to change a word throughout your file. The common format for this is

1,$ s/*pattern*/ *replacement*/g

This tells `ex` to make the replacement on every line of the file (1,$ means from the first line to the last line); `g` tells it to make the replacement every time it appears in a line. If you didn't add the `g`, it would make the replacement only once per line, at the first appearance of *pattern*.

A variation of this format is useful when you want to repeat the same change several times within one line. Instead of giving 1,$ for the range of the command, you may specify a single line. For example, if the above sample line were the current line (or dot) in your file, you could type

```
s/e/a/g
```

to change your line to

```
Sand Tha Claaning Bill To My Gattysburg Addrass
```

In general, `substitute` replaces the first instance of *pattern* with *replacement* on each specified line. The *suffixes* are

g   (global.) Substitute *pattern* with *replacement* every time it appears on the specified lines. To make the substitution everywhere in the file, use the format

           1,$s/*pattern*/ *replacement*/g

           You can also use % instead of 1,$.

c   (confirm.) Print the line before making each substitution, marking the string to substitute with ^ characters. Type y to confirm the substitution, and type any other character if you don't want to make the substitution.

r   (replace.) Replace the previous replacement pattern from a substitution with the most recent search string.

You can split lines by substituting newline characters into them. You must escape the newline in *replacement* by preceding it with a backslash (\). See "Regular Expressions and Searching" for other metacharacters available in *pattern* and *replacement*.

Omitting *pattern* and *replacement* repeats the last substitution. For example, if you substitute the word `test2` for the word `test` on one line using the command

```
s/test/test2/
```

you can repeat this substitution on another line by typing

```
s
```

This is a synonym for the `&` command, which is described later in this chapter.

Using the `r` suffix (`sr`) replaces the previous *pattern* with the previous regular expression. For example, if you make the substitution

```
s/test/test2
```

then search for a *pattern* such as `my.test`,

```
/my.test/
```

then the command

```
sr
```

changes `my.test` to `test2`. If you omit the `r` suffix, the `s` command replaces `my.test` with `my.test2`. This is a synonym for the `~` command, which is described later in this chapter. See "Command Summary" and "Repeating the Last Command" in Chapter 3, "Using the `vi` Screen Editor," for more information on repeating substitutions.

Copying and moving text

You can use several commands to copy and move text.

The `copy` command places a copy of one section of your text after the specified line. The most common format of this command is

line1, *line2*`co[py]` *lineno*

This copies the lines between *line1* and *line2* and places them after *lineno*.

The `move` command moves a section of your text to a new location in your file. The format of this command is

[*line1*[, *line2*]]m[ove] *lineno*

This moves the text either from *line1* or between *line1* and *line2* after *lineno*.

You can also move text by moving it into `ex`'s buffer and then placing it in the text with the `put` command.

There are two general forms of commands used for moving text in this way. The first is to use the `yank` or `delete` commands to place the text in `ex`'s unnamed buffer. This happens automatically when you use the `delete` command to delete the text (described under "Deleting Text"), or the `yank` command to copy the text. The general forms for the `yank` command are

[*line1*[, *line2*]]ya[nk]

to place a copy of the text either from *line1* or between *line1* and *line2* in `ex`'s unnamed buffer, or

ya[nk] *n*

to place a copy of the next *n* lines in `ex`'s unnamed buffer. You can then place `ex`'s buffer somewhere else in the file with the `put` command (either by moving to where you want the text to appear, or by specifying an address before the `put` command). If you use `ex`'s unnamed buffer, you can't make any modifications to your text between placing the text in your buffer and putting it in its new location.

The other way to move or copy text is to use one of `ex`'s named buffers. `ex` has 26 buffers named `a` through `z`. Use the same general format as before, but specify a buffer name with each command. For example,

1,4 d a

deletes lines 1 through 4 and places them in buffer `a`.

10 pu a

puts the contents of buffer `a` after line 10.

You can also specify the buffers as `A` through `Z` if you want the text you are currently deleting or yanking to be appended to the end of the buffer rather than overwriting it.

You can use `ex`'s named buffers to move information from one file to another if you have specified both files when you started `ex`.

Regular expressions and searching

A regular expression uses **metacharacters** (special characters that stand for other characters) to stand for a set of strings. The regular expression is said to match each element in this set of strings. For example, if . is a special character standing for any letter and A is an ordinary character standing for A, then A. would find all of the following words: At, About, Another.

Regular expressions in ex always appear between the characters / / or the characters ? ?.

char	Any character other than the metacharacters listed below matches itself. The characters listed below are metacharacters. You have to precede them with a backslash (\) to have ex treat them as ordinary characters.
^	At the beginning of a pattern, the caret specifies that the pattern is at the beginning of a line. For example, ^A specifies a line beginning with A. This character has a different meaning within square brackets.
$	At the end of a pattern, the dollar sign specifies that the pattern is at the end of a line. For example, a$ specifies a line ending with a.
.	The period matches any single character except the newline character. For example, A. matches A followed by any character.
[*pattern*]	A pattern enclosed in square brackets sequentially matches a set of single characters defined by *pattern*.
	Ordinary characters in *pattern* match themselves. For example, [ab] specifies either a or b.
	A pair of ordinary characters separated by - in *pattern* defines a range of characters. For example [a-z] matches any lowercase letter.
	If ^ is the first character within square brackets, it specifies characters that are not in the pattern. For example, [^a-z] matches anything but a lowercase letter.
\<	This matches a pattern at the start of a word (ex defines the start of a word as the beginning of a line, or a letter, digit, or underline that follows any character other than a letter, digit, or underline).
\>	This matches a pattern at the end of a word (ex defines the end of a word as the end of a line, or a letter, digit, or underline followed by any character other than a letter, digit, or underline).

You can use the preceding regular expressions to construct larger regular expressions using the following rules:

- If you use two regular expressions (for example, `[a-z][A-Z]`), the editor matches the first string it encounters that matches both regular expressions in the order they appear.

- Following a regular expression with an asterisk (`*`) matches *zero* or *more* occurrences of the preceding character. Generally, you should use this within a longer regular expression, since it matches for zero occurrences first. That is, searching for `a*` finds zero occurrences and matches the characters following the cursor. This is convenient, however, for longer regular expressions. For example, `ba*b` matches `bb`, `bab`, `baab`, `baaab`, and so on. Within a longer regular expression, if there is any choice, it matches the longest leftmost string. In the preceding example, it would choose `baaab`.

- Enclosing a regular expression in `\(` and `\)` defines the regular expression as a numbered "field," so that you can refer to it later. For example,

  `\([a-z]\)\([A-Z]\)`

  defines two fields, numbered sequentially. Use `\`*f* (where *f* is the number of the field) to refer to any of the fields in the regular expression. In the previous example `\1` refers to the field matched by the regular expression `[a-z]` and `\2` refers to the field matched by the regular expression `[A-Z]`.

- The null regular expression (`//` or `??`) is shorthand for the last regular expression.

Turning off metacharacters

There are two ways to use `ex`'s metacharacters as ordinary characters (if, for example, you want to search for the period [`.`] character).

1. You can precede the metacharacter with a backslash (`\`).

2. You can set the `nomagic` option (see "Setting Options"). This strips all but the following three metacharacters of their special meaning: `^` at the beginning of a regular expression (indicating the beginning of a line), `$` at the end of a regular expression (indicating the end of a line), and the backslash character (`\`). You can restore the special meaning to the other metacharacters by preceding them with a backslash (`\`).

Working with multiple files

You can work with several files during one editing session. Before describing the commands for this, we define the terms that refer to these files.

Working with the current file

The file you are editing is the **current file.** This means that the buffer contains the edited version of this file. `ex` overwrites this file with the updated version in the buffer without protest. When `ex`'s current file is not the file being modified in the buffer (for example, if you use the `file` command to change the name of the current file without changing the buffer).

`f` *filename*

`ex` does not overwrite a file with the buffer's contents. See "Changing the Current File." If the file in the buffer is not the current file, the `file` command displays

```
[Not Edited]
```

You can use `%` to refer to the current file anywhere you would use the filename.

Working with alternate files

`ex` also has an **alternate file,** which is usually the previous file you edited.

If you haven't previously edited another file, but have read a file into the buffer with the `read` command, this becomes the alternate file.

You can use `#` to refer to the alternate file anywhere you would use the filename. This makes it easy to alternate between two files. For example, if you are in the current file and want to edit the alternate file, type

```
e #
```

This reads the alternate file into the buffer, making it the current file. (If you had modified the buffer since you last wrote it to file, `ex` would warn you and would not edit the alternate file.)

Opening multiple files at startup

You can specify more than one file when you start `ex`. The format is

`ex` *filename* `...`

`ex` reads the first file into the buffer and creates an **argument list,** which is a list containing the names of all the files you specified at startup.

Displaying the argument list

The `args` command displays the current argument list with the current file delimited by brackets (`[ ]`).

Editing the next file on the argument list

The `next` command edits the next file in the argument list. The format for this command is

`n[ext]`

You must save any changes you have made before editing a new file or you'll get the message

`No write since last change (:next! overrides)`

To edit the next file in the argument list and overwrite the current buffer with this file, type

`n!`

The buffer is overwritten with the next file, whether you've saved the current buffer or not.

Typing

`n+`*cmd*

runs *cmd* after opening the first file on the argument list.

Replacing the argument list

You can replace the list of files in the argument list with another list of files by typing

n *filename* . . .

`ex` replaces the list of files in the argument list and edits the first file on the new list.

If you made changes to one or more files on the original argument list and you haven't saved the current buffer, you'll get the message

```
No write since last change (:n! overrides)
```

If you haven't saved the buffer, you can force `ex` to replace the argument list with the new list by typing

n! *filename* . . .

The new list of files replaces the original one even if you have not saved the current buffer.

Returning to the first file on the argument list

The *rewind* command edits the files on the argument list beginning with the first file. The format of the command is

```
rew
```

If you made changes to one or more files on the original argument list and you haven't saved the current buffer, you'll get the message

```
No write since last change (:rewind! overrides)
```

Typing

```
rew!
```

forces `ex` to edit the files on the argument list beginning with the first file and discarding any changes you made to the current buffer.

Editing a new file

The `edit` command reads a new file into the buffer. If you haven't saved the current buffer, `ex` warns you and doesn't edit the new file. If you have saved the current buffer, `ex` deletes the buffer contents, makes the specified file the current file, and prints the new filename and the number of lines and characters read.

    `ex` sets the current line to the last line in the new file (in line mode) or the first line in the new file (in open or visual mode).

    `ex` strips the high-order bit from any non-ASCII characters and discards any null characters.

    If the file is a special device (block, character, or TTY), `ex` tells you and allows you to edit the file. If it is a directory, you'll get the message `Directory`. If it is a binary file, you'll get the message `Line too long` or `Incomplete last line`. Generally, you shouldn't edit these kinds of files.

    The `edit!` command edits the specified file and overwrites the current buffer, whether you've written it or not. The form of this command is

`e!` *filename*

Typing

`e+` *lineno filename*

or

`e+/` *pattern filename*

begins editing *filename* at line *lineno* or at pattern */pattern* (*pattern* can't contain spaces).

Copying another file to the current buffer

The `read` command copies the text of *filename* to the current buffer after the specified line. The format is

[*lineno*]`r[ead]` [*filename*]

where *lineno* is a line number, or an expression resulting in one. If you don't specify a *filename*, it uses the current filename. If there is no current filename, *filename* becomes the current name. If the file buffer is empty and there is no current file, `ex` treats this as an `edit` command.

read then tells you the filename read in and the number of lines and characters read. After a read command, the current line is the last line read (in ex) or the first line read (in open or visual mode).

Typing

```
0read
```

reads the file at the beginning of the buffer.

Examining the characteristics of the current file

The file command tells about the file you are editing. Its syntax is

```
f[ile]
```

It prints the following information:

- the current filename
- whether you have modified the current file since the last write
- whether the current file is in read-only mode
- the current line
- the number of lines in the buffer
- the current line's position in the buffer, given as a percentage from the beginning of the buffer

It also notes when the current file is "not edited" (the current file is not the file in the buffer). In this case, you have to use w! to write to the file, since ex does not know if write will destroy a file unrelated to the current buffer contents.

Changing the current file

The file command changes the current filename to *filename* without changing the buffer contents. The format of the command is

```
f[ile] filename
```

The current file is then considered "not edited."

Using shell commands in `ex`

There are several ways of interacting with the shell from within `ex`.

Running another program from `ex`

The `!` character invokes a shell to run a single command using the syntax `!` *command.*
This runs *command* in the shell and returns you to the editor when the command
completes.

The command you invoke from the editor using the `!` syntax may be a simple
command such as `ls`, or it may be an interactive program such as `dc` or a shell script.

If you enter a simple shell command from the editor, it runs immediately and prints
`!` on the screen when it ends. You are then back in the editor at the same position in
the file. For example,

```
!ls
```

lists the files in your current directory. If you haven't written the buffer contents since the
last change, `ex` prints a warning message before executing the command. Before
returning you to the editor, it prints "`!`".

If you enter an interactive program from the editor, it runs until you exit that
program. For example,

```
!dc
```

invokes the `dc` calculator program. You can then use `dc` as long as you wish. When
you end `dc`, you are back in the editor at the same position in the file.

The command

```
sh
```

invokes your login shell. You may then give as many commands in the shell as you wish.
When you finish with the shell, type an *eof* character (CONTROL-D in the A/UX standard
distribution) or

```
exit
```

to return to the editor.

You may escape to a shell different from your login shell. The general form for this command is

!shell

where *shell* is the name of the shell you wish to invoke. For instance, if your login shell is the Bourne shell (`sh(1)`), you may invoke the `csh` instead with the following:

`!csh`

This invokes a copy of the C shell, temporarily suspending `ex`. You may then give as many commands in the new shell as you wish. When you finish with the shell, type an *eof* character (Control-D in A/UX) or

`exit`

to return to the editor.

Remember that after you use the `sh` (or *!shell*) escape from `ex`, you have invoked a new shell, not the shell from which you initiated `ex`. If you use `sh` to escape `ex` and forget that you have suspended your editing job, you might invoke a new copy of `ex` from your new shell instead of exiting that shell and going back to the original `ex` session. This can cause problems with inconsistent versions of a file if you finally quit `ex`, and is something to be aware of when using `sh` or *!shell* from `ex`.

Directing command output to the buffer

You can read the output of a command into your file with the `read !` command. The usual format of this command is `r[ead] !`*command*
(you must type a space or a tab before the `!`). For example,

`read !ls`

places the directory listing after the current line.

◆ **Note** The shell prompts are also written to the file. ◆

Sending the buffer to shell commands

You can send part of your buffer to a command with the `write !` command (you must precede the `!` with a space or a tab). The format of this command is

line2[, *line2*]]w[rite] `!` *command*

For example, to format the first 20 lines of your file without leaving the editor, use the command

```
1,20 w !nroff > new.file
```

When you precede this command with a range of line numbers, `ex` sends the specified line(s) to *command* and replaces the line(s) with the output of the command.

Writing shell scripts using `ex` commands

You can write shell scripts that use `ex` commands. You add comments to these scripts by starting a line with a double quote (`"`) or by adding a double quote and a comment to the end of a command (except when they could be read as part of the command—as in shell escapes and the `substitute` and `map` commands).

You can place more than one command on a line by separating each pair of commands with a `|` (pipe) character. If you use a global command, comment, or shell escape (`!`), however, it must be the last command on the line.

You can use the `—` flag option to `ex` within shell scripts to suppress interactive feedback. This permits the script to run without pausing for information to be typed in from the terminal.

The following is an example of a shell script:

```
for i in $*
do
ex - $i <<EOF
g/f1/s//fR/g "change \f1 to \fR
g/f2/s//fI/g "change \f2 to \fI
g/f3/s//fB/g "change \f3 to \fB
wq
EOF
done
```

To run this shell script (named `script.ex`) on a text file, make the script file executable with the command

```
chmod +x script.ex
```

Then type

```
script.ex filename
```

where *filename* is a text file. When the script has finished running, all instances of "`\f1`" in your file are changed to "`\fR`," and so on. See also Chapter 6, "Using the `sed` Stream Editor," for information on making global changes to a file using scripts.

Setting options

You control many of the ways `ex` behaves by setting options.

`ex` has three kinds of options: numeric, string, and toggle. Each of these options is set in its own way. **Numeric options** are options that take a numeric value and **string options** are options that take a string value. You set numeric and string options with the following command format:

```
set opt=val
```

For example, you set the number of lines to scroll through (a numeric option) with the following command:

```
set scroll=4
```

You set the default shell (a string option), with the following command:

```
set sh=/bin/sh
```

A **toggle option** is an option that is either on or off. You set a toggle option with one of the following formats:

```
set opt
```

turns the option on and

```
set noopt
```

turns the option off. For example,

```
set number
```

turns on line numbering and

```
set nonumber
```

turns off line numbering.

To set options, use the following syntax:

Option type	Syntax	Sets opt to
numeric	set *opt=number*	*number*
string	set *opt=string*	*string*
toggle	set *opt* on	
	set no*opt* off	

You can place multiple options on one line with the format

```
set opt opt opt
```

Most options can be abbreviated; see "Summary of Command Options," later in this chapter, for a list of options and their abbreviations.

Listing options

To see option settings, use the following syntax:

Options listed	Syntax
all	set all
changed ones	set
an *opt*	set *opt*?

When to set options

You can set these options anytime while editing a file, or you can set them as part of your default editing environment by including them in the `EXINIT` variable in your `.profile` file or by creating a `.exrc` file in the current or home directory. If you set them in your `.profile` file, they should all be on one line, separated by vertical bar (`|`) characters. For example, in your `.profile` file, you could have the following entry:

```
EXINIT="set number|set scroll=20|set terse"
```

Command option summary and descriptions

This section provides a brief description of the function of each of the command options for the `ex` line editor.

autoindent

Abbreviation: `ai`

Default: `noautoindent`

Begin new lines of text at the indent level of the previous line. This is useful in structured program text.

When inserting text, CONTROL-D moves the cursor back to the previous tab stop.

Entering a blank line or typing `0`CONTROL-D erases the autoindent. You can type one line at the margin by beginning it with `^`CONTROL-D; the next line returns to the previous indent.

`autoindent` does not work with global commands or when the input device is not a terminal.

autoprint

Abbreviation: `ap`

Default: `autoprint`

Print the current line after each `delete`, `copy` (or `t`), `join`, `move`, `s` (substitute), `undo`, `<`, or `>` command (if it is the last command on the line). This is suppressed during global commands.

autowrite

Abbreviation: aw

Default: noautowrite

Write the modified buffer contents to the current file when you use the following ex commands

n[ext]	Edit next file in series
rew[ind]	Reedit the list of files from the start
ta[g]	Move to a tag location
!	Escape to the shell

or the following vi commands:

CONTROL-^	Switch files
CONTROL-]	Move to a tag location

You can override the autowrite option, destroying the current buffer contents, with the following ex commands

e[dit]	Instead of n[ext]
rew[ind]!	Instead of rew[ind]
ta[g]!	Instead of ta[g]
sh[ell]	Instead of !

and the following vi commands:

:e #	Instead of CONTROL-^
:ta!	Instead of CONTROL-]

beautify

Abbreviation: bf

Default: nobeautify

Discard all control characters (except tab, newline, and form feed) from your input, and print a message the first time it discards a backspace character. This option applies only to text input and not to command input.

directory=*dir*

Abbreviation: dir

Default: directory=/tmp

Specify the directory where ex places its buffer file. This directory must have write privileges, or the ex application will quit.

edcompatible

Abbreviation: none

Default: noedcompatible

Use the suffixes g and c to toggle globally and confirm settings of the s (substitute) command.

errorbells

Abbreviation: eb

Default: noerrorbells

Sound a bell when displaying error messages (you cannot suppress this bell in open or visual mode). If possible, ex highlights the error message on the screen instead of ringing the bell.

flash

Abbreviation: fl

Default: flash

Flash the screen when an error occurs.

hardtabs=*n*

Abbreviation: ht

Default: hardtabs=8

Set the length of terminal hardware tabs (or where the system expands tabs).

ignorecase

Abbreviation: ic

Default: noignorecase

Set regular expressions to match both uppercase and lowercase patterns, except when you specify a range of uppercase characters (for example, [A-Z]).

insertarrows

Abbreviation: ia

Default: insertarrows

Allow use of arrow keys in insert mode as well as command mode.

lisp

Abbreviation: none

Default: `nolisp`

Set `autoindent` and modify the `vi` motion commands `()`, `{}`, `[[`, and `]]` to have meaning for LISP.

list

Abbreviation: `list`

Default: `nolist`

Print lines showing tabs as `^I` and the end of the line as `$`, as in the `list` command.

magic

Abbreviation: none

Default: `magic`

With `magic` set, `ex` recognizes all the metacharacters used in regular expressions. Setting `nomagic` uses only the following regular-expression metacharacters: `\`, `^`, and `$`. It treats all other characters (including `~` and `&` used in substitutions) as normal characters. You can use any of these as metacharacters by preceding them with `\`.

mesg

Abbreviation: none

Default: `mesg`

By default, `vi` allows other users to send you messages while you are editing a file. `nomesg` turns off this permission.

number

Abbreviation: `nu`

Default: `nonumber`

Print lines preceded by their line number and (after RETURN) prompt with line numbers for input lines.

open

Abbreviation: none

Default: open.

By default, you can enter open and visual mode from ex. Setting noopen means you can't use these modes.

optimize

Abbreviation: opt

Default: optimize

Suppress carriage returns on more than one (logical) line of output. This optimizes output on terminals without addressable cursors when printing text with leading white space.

paragraphs=*xyz*

Abbreviation: para

Default: paragraphs=PLIbp

Specify the paragraph macro searched for in vi and open mode when you type { or }. By default, it searches for the mm macros .P, .LI, and the nroff/troff request .bp.

prompt

Abbreviation: none

Default: prompt

By default, ex prints the prompt (:) when it is in command mode. Setting noprompt turns off this prompt.

readonly

Abbreviation: none

Default: noreadonly

Make the file read only (just as if you had started ex with the -r flag option). You can override this option and save the file by using the write! command.

redraw

Abbreviation: none

Default: `noredraw`

Force a dumb terminal to redraw the characters to the right of the cursor as you type input in `vi`. This is useful only at baud rates of 1200 or higher.

remap

Abbreviation: none

Default: `remap`

Repeatedly translate maps until they are unchanged. For example, if you map o to O, and O to I, setting `remap` maps o to I, while setting `noremap` maps o to O.

report=*n*

Abbreviation: none

Default: `report=5`

Print a message when a command modifies more than *n* lines. For example, `ex` prints `12 lines deleted` after a deletion. For the following commands, `ex` reports after completing the entire command: `global, open, undo,` and `visual.`

scroll=*n*

Abbreviation: none

Default: `scroll=`half the value of the `window` option

Set how many lines scroll when you press CONTROL-D in `vi`'s command mode. By default, it uses half the number of lines set with the `window` option.

sections=*xyz*

Abbreviation: none

Default: `sections=HHU`

Specify the section macro searched for in `visual` and `open` mode when you type `[[` or `]]`. By default, it searches for mm's, `.H`, and `.HU` macros.

shell=*pathname*

Abbreviation: `sh`

Default: `shell=$SHELL`

Set the pathname of the shell used by the shell escape command `!` and the `shell` command. `$SHELL` is the value in the `SHELL` variable, as set in the `.login` or `.profile` file.

shiftwidth=*n*

Abbreviation: `sw`

Default: `shiftwidth=8`

Set the software tab stop width used when you press CONTROL-D in `autoindent`, and when you use the `>` and `<` commands.

showmatch

Abbreviation: `sm`

Default: `noshowmatch`

Move the cursor to the matching `(` or `{` on the screen for one second when you type `(` or `{` in `vi`. Extremely useful with LISP.

slowopen

Abbreviation: `slow`

Default: `noslow`

Don't update the display when you enter text in `vi`. Useful on a very slow line.

tabstop=*n*

Abbreviation: `ts`

Default: `tabstop=8`

Set tabstops to *n*.

taglength=*n*

Abbreviation: tl

Default: taglength=0

Tags are not significant beyond *n* characters. Zero (the default) means that all characters are significant.

tags=*pathname*

Abbreviation: none

Default: tags=/usr/lib/tags

Search for the requested files sequentially in the specified path when using the tag command. You must escape any spaces with a backslash (\). By default, it searches in the current directory and in /usr/lib (a master file for the entire system).

term=*terminal*

Abbreviation: none

Default: $TERM

Set your terminal type. The value you specify must exist in a file in the appropriate subdirectory of /usr/lib/terminfo. $TERM is the value of the TERM variable, as set in the .login or .profile file.

terse

Abbreviation: none

Default: noterse

Produce shorter error messages.

warn

Abbreviation: none

Default: warn

Print [No write since last change] if you use a ! *command* escape before you have saved the current buffer.

window=*n*

Abbreviation: none

Default: speed dependent

Set the number of lines in `vi`'s text window. By default, this is determined by your baud rate. The default settings are

8   for 600 baud or lower
16   for 1200 baud
23   (or full screen) at higher speeds

The settings w300, w1200, w9600 set `window` only if the speed is slow (300), medium (1200), or high (9600), respectively. They are suitable for an `EXINIT` variable.

wrapmargin=*n*

Abbreviation: wm

Default: wrapmargin=0

Set the column number where the cursor automatically returns when you enter text in visual and open modes. This is determined by setting the wraparound point *n* columns from the right side of the screen. Since there are 80 columns on a typical terminal screen, setting `wrapmargin=10` would break a long line 10 columns to the right of this, in the 70th column. `wrapmargin=0` means that a long line wraps at column 80, but is not broken (that is, `wrapmargin` is off).

wrapscan

Abbreviation: ws

Default: wrapscan

Search the entire file for a regular expression by moving from the current line, wrapping around the end or beginning of the file (depending on the direction you're searching in), and returning to the current line.

writeany

Abbreviation: wa

Default: nowriteany

By default, `ex` warns you if you try to save your buffer to any file other than the current file. Setting `nowriteany` allows you to write to any file with write permission.

Summary of `ex` options

Table 4-3 is a complete list of `ex` editor command options.

Table 4-3 Summary of `ex` options

Option	Abbreviation	Default
autoindent	ai	noai
autoprint	ap	ap
autowrite	aw	noaw
beautify	bf	nobf
directory	dir	dir=/tmp
edcompatible	-	noedcompatible
errorbells	eb	noeb
hardtabs	ht	ht=8
ignorecase	ic	noic
lisp	-	nolisp
list	-	nolist
magic	-	magic
mesg	..	mesg
number	nu	nonu
open	-	open
optimize	opt	opt
paragraphs	para	para=PLIbp
prompt		prompt
readonly	-	noreadonly
redraw	-	noredraw
remap	..	remap
report	-	report=5
scroll	-	scroll=1/2 window
sections	-	sections=HHU
shell	sh	sh=$SHELL

Table 4-3 Summary of ex options (*continued*)

Option	Abbreviation	Default
shiftwidth	sw	sw=8
showmatch	sm	nosm
slowopen	slow	(terminal dependent)
tabstop	ts	ts=8
taglength	tl	tl=0
tags	–	tags=/usr/lib/tags
term	–	term=$TERM
terse	–	noterse
warn	–	warn
window	–	window= (speed dependent)
wrapmargin	wm	wm=0
wrapscan	ws	ws
writeany	wa	nowa

Mapping and abbreviations

Mapping and abbreviations are available as a joint facility of ex and vi. The mapping or abbreviation must be defined on the ex command line, but is useful only in visual mode. See "Mapping and Abbreviations" in Chapter 3, "Using the vi Screen Editor," for further information.

Additional ex commands

This section describes several other ex commands. See "Command Summary" for a complete list of ex commands and their usage.

Marking text

You can set up a special address with the `mark` command and then use this address anywhere you would use an `ex` address. The syntax for this command is

[*lineno*]ma[rk] *x*

where *lineno* is the line (number) in the file to mark and *x* is a lowercase letter to mark this line. You must precede this letter with a blank or a tab. After you have marked a line, you can refer to it by typing

'*x*

where *x* is the name you gave it.

You can also use the `k` command to mark a line. The syntax for this command is

[*lineno*]k*x*

where *lineno* is the line (number) in the file to mark and *x* is a lowercase letter to mark this line. This is the same as the `mark` command, except you don't have to precede the letter with a space.

Recovering lost text

There are several ways to recover information in `ex`.

The `undo` command changes the buffer back to the way it was before the last editing command.

If you have deleted several large sections and want to recover them, the easiest method is to exit `ex` without saving your changes by using the `quit!` command. The format for this command is

q[uit]!

This restores the buffer you had when you last wrote the file.

The `preserve` command is a more drastic way to save your file. You should use it only if you can't save your file with a `write` command. It saves your file as though your system had just crashed. You can recover this file with the `recover` command.

Editing programs

`ex` has several commands for editing programs.

`<` shifts text right to the next tabstop and `>` shifts text left. This is useful for changing the indentation of a section of program.

Using the `tag` command helps you locate functions that may be spread over many files. The `tag` command starts editing the file at *tag,* moving to another file if necessary. Since the current file edit may be canceled abruptly, you must write the current file, if you modified it, before giving a `tag` command. If you give the `tag` command without specifying a *tag,* `ex` uses the previous tag. The syntax is

`ta[g]` *tag*

Normally, you use this command after using the `ctags` command to create a tag file (see `ctags`(1) in *A/UX Command Reference*). A tag file consists of several lines with three fields separated by blanks or tabs. The first field is the name of the tag, the second is the name of the file where the tag resides, and the third is an addressing form used to find the tag. Usually, this is a contextual scan in the form */pattern/,* performed as if `nomagic` were set.

Names in the tag file must be sorted alphabetically.

For instance, if you wish to have ready access to the functions in a file called `functions.c` of the following form,

```
main ()
{
        f1 ();
        f2 ();
}
f1 ();
{
        f2 ();
}
f2 ();
{
}
```

you could run the `ctags` program on it, as follows:

```
ctags functions.c
```

This creates a tag file called `tags`.

If you then edit `functions.c` *or even another file* with `ex`, and give the command

```
ta f1
```

you will find yourself in the editor buffer of `functions.c` at the line

```
f1()
```

ready to edit that function. However, you must be in the same directory as the `tags` file.

Saving text and quitting `ex`

The `write` command saves the changes you've made to the buffer in the specified file and prints the number of lines and characters written. The format of this command is

[*line1*[, *line2*]]w[rite] [*filename*]

By default, `ex` writes the entire buffer to *filename*. Including a starting and ending address saves only the lines between *line1* and *line2*. Including just a starting address saves only *line1*. If you don't include *filename*, `ex` writes to the current file. If there is no current file, `ex` creates the file *filename* and writes the buffer to it. The `write` command also writes to `/dev/tty` and `/dev/null`.

The `write!` command forces `ex` to write the buffer. The file must already exist. The format of this command is

[*line1*[, *line2*]]w[rite!] [*filename*]

The `write>>` *filename* command appends the buffer to the end of a file (which must already exist). The format of this command is

[*line1*[, *line2*]]w[rite] >> *filename*

Quitting ex

The `quit` command cancels `ex` using the syntax.

q[uit] RETURN

If you haven't written all your changes to a file, `ex` warns you and does not quit. It also tells you if there are additional files in the argument list.

The `quit!` command allows you to leave `ex` without saving the changes you've made, using the syntax

q[uit]! RETURN

See "Saving Files and Quitting `ex`," earlier in this chapter, for information on saving a file.

Saving a file and quitting ex simultaneously

The following commands save your changes and quit the `ex` application.

The `write` and `quit` commands can be used together to save your changes in the current file or in a specified file and then quit `ex`.

wq

This is simply a shorthand form of using `write` followed by `quit`; you may also use the long form of these commands as shown in the sections above.

Followed by an exclamation point,

wq! *filename*

forces `ex` to save your changes in filename and then exits `ex`. This is a shorthand form of using `write!` followed by `quit`. You may also use the long form of these commands as shown in the sections above.

The `xit` command saves your buffer only if you have made changes since last saving it and then exits `ex`. You can use this command by typing

x

The long format of this command is

x[it] [*filename*]

where *filename* is the name of the file in which to save your buffer.

Error conditions

When there is an error, the system beeps. If `ex` receives an interrupt signal, it prints `Interrupt` and returns to command mode. If the primary input is from a file, `ex` quits.

Limitations

The `ex` line editor cannot exceed the following parameters:

- 1024 characters per line
- 256 characters per global command list
- 128 characters per filename
- 128 characters in the previous inserted and deleted text in open or visual mode
- 100 characters in a shell escape command
- 63 characters in a string option
- 30 characters in a tag name

A limit of 250,000 lines in the file is silently enforced.

The number of macros defined with `map` in `vi` is limited to 32, and the total number of characters in macros is limited to fewer than 512.

Recovering lost files

If the system crashes or you accidentally hang up before saving your file, `ex` sends you a message informing you that it has saved your file. The `-r` flag option recovers your buffer.

The format of this command is

```
ex -r
```

to list the files that `ex` saved, and

```
ex -r filename
```

to recover *filename* (after first moving to the directory you were in). You should check the recovered file before overwriting the existing file with it. Note that you must use the write command to explicitly write the recovered file. The `xit` command is a safe way to exit a recovered file because it does not guarantee that a write operation will occur.

Command summary

In the following commands, the last line you enter, copy, change, or print becomes the current line. If you use an input command, but don't enter a new line, the next line becomes the current line. If you delete text, the following line become the current line; deleting text at the end of the file makes the new last line the current line.

The following is the standard `ex` command format:

[*line-spec*] *command*[!] [*parameters*] [*n*] [*flags*]

Only pressing RETURN prints the next line. The `ex` editor ignores a period character (.) preceding any command.

In these commands,

line-spec	(for "line-specifier") indicates the command address (see "Line Selection" for the definition of *line-spec*). In the following commands, *line-spec* defaults to the current line, unless stated otherwise.
lineno	(for "line number") indicates the single-line command address (defaulting to the current line).
flags	indicates invoking the printing command `#` (number), `l` (list), or `p` (print) after the command.
command n	performs *command* on the *n* lines involved, starting at the current line.

ab[breviate] *wd word*

(Must be defined in `ex`, but works in `vi` only.) Abbreviate *word* to *wd* in input mode. Typing *wd* in input mode, delimited by spaces or punctuation, displays *word*.

[*lineno*] a [ppend] RETURN *text* RETURN.

Append *text* after *lineno*. Specify *lineno* zero (0) to insert text at the beginning of the buffer. Default address: current line.

[*lineno*] a [ppend] ! RETURN *text* RETURN.

Same as `append`, but changes the setting for `autoindent` while appending (see "Setting Options"). Default address: current line.

ar[gs]

Print the list of files to edit, with the current file delimited by brackets (`[ ]`).

[*line1* [, *line2*]] c [hange] [*n*] RETURN *text* RETURN.

Replace *line1*, or the lines between *line1* and *line2*, or *n* lines, with *text*. Default address: current line.

[*line1* [, *line2*]] c [hange] ! [*n*] RETURN *text* RETURN.

Changes the setting of the `autoindent` option while changing the text (see "Setting Options"). Default address: current line.

[*line1* [, *line2*]] co [py] *lineno* *[flags]*

Copy the text on *line1*, or between *line1* and *line2*, after *lineno* (if *lineno* is zero (0), the text is copied to the start of the file). Including a *flag* after the command (either `p` for `print`, `l` for `list`, or `#` for `number`) changes the display to the specified format. `t` is a synonym for `copy`. Default address: current line.

[*line1* [, *line2*]] d [elete] *[buffer]* *[n]* *[flags]*

Delete the specified lines (either *line1*, those lines between *line1* and *line2*, or the number of lines specified with *n*). Specifying *buffer* with a lowercase letter (`a` through `z`), overwrites that buffer with the deleted text, while specifying an uppercase letter (`A` through `Z`) appends the deleted text to that buffer. Including a *flag* after the command (either `p` for `print`, `l` for `list`, or `#` for `number`) changes the display to the specified format. Default address: current line.

e[dit] *file*

Edit *file*. `ex` warns you if you haven't written your current buffer, and doesn't allow you to edit *file*. Otherwise, `ex` reads *file* into the buffer, making it the current file, and prints the new filename and the number of lines and characters read. `ex` strips the high-order bit from any non-ASCII characters and discards any null characters.

The current line is the last line (in `ex`), or the first line (in `open` or `visual` mode).

If *file* is a special device (block, character, or TTY), `ex` mentions this, but allows you to edit the file. If it is not a text file, it prints the error message `Line too long` (or `Directory`, if appropriate).

e[dit]! *file*

Same as `edit`, but `ex` doesn't warn you if you haven't saved the current buffer.

e[dit]+lineno** *file*

Begin editing *file* at line *lineno*, where *lineno* may be a line number or a pattern /*pattern* (*pattern* can't contain spaces).

f[ile]

Print the following information: current filename; whether you have saved the current buffer; whether the current file is in read-only mode; the current line; the number of lines in the buffer; the percentage of the way through the buffer of the current line. It also notes when the current file is not the file in the buffer, by printing `not edited`. This happens if you changed the current file by entering `file` *file*. In this case, you have to use `w!` to write to the current file.

f[ile] *file*

Change the current filename to *file*, without changing the buffer. *file* is considered "not edited."

[line1[,line2]]g[lobal]/pat[/cmds/]

`global`/*pat* prints lines containing the regular expression pat and `global`/*pat*/*cmds*/ performs *cmds* at lines containing the regular expression *pat*.

cmds can span several lines if you end all but the last line with a backslash (\). *cmds* can include `append`, `insert`, or `change`. If one of these is the last command, you can omit the period that cancels these commands. *cmds* can also include the `open` or

visual commands, which take input from the terminal. *cmds* cannot include global or undo, since undo would reverse the entire global command.

global turns off the autoprint and autoindent options and sets the report option to infinity until executing the entire command. ex sets the context mark (") to the current line, and changes it only if you enter open or visual mode within the global command. Default address: 1,$.

[*line1* [*, line2*]] g [lobal] ! /*pat*/*cmds* or [*line1* [*, line2*]] v /*pat*/*cmds*

Run *cmds* globally on each line *not* matching *pat*.

[*lineno*] i [nsert] Return [*text* Return] .

Insert text before *lineno*. This command differs from append only in its placement of text. Default address: current line.

[*lineno*] i [nsert] ! Return [*text* Return] .

Same as insert, but changes the autoindent option while inserting (see "Setting Options"). Default address: current line.

[*line1* [*, line2*]] j [oin] [*n*] [*flags*]

Join the specified lines (either the lines between *line1* and *line2* or *n* lines). ex ensures that there is at least one blank character where the lines joined, two if there was a period at the end of the line, or none if the first following character is a). If there is already white space at the end of the line, it discards the white space at the start of the next line. Including a *flag* after the command (either p for print, l for list, or # for number) changes the display to the specified format. Default address: current line.

[*line1* [*, line2*]] j [oin] ! [*n*] [*flags*]

Same as join, but doesn't add or delete white space. Default address: current line.

kx

A synonym for mark (described below), which does not require a blank or tab before the letter. Default address: current line.

[*line1* [*, line2*]] l [ist] [*n*] [*flags*]

Print the specified line(s), displaying tabs as ^I and the end of line(s) as $. Including a *flag* after the command (either p for print or # for number) changes the display to the specified format. Default address: current line.

map *string definition*

(Must be set via an `ex` command, but works in `vi` only.) Make typing string equivalent to typing definition when executing commands in visual mode (`vi`). String can be up to 10 characters long, and definition can be up to 100 characters long. Default address: none.

[lineno] **ma[rk]** *x*

Mark the specified line with *x*, a lowercase letter. You must precede *x* with a blank or a tab. After marking a line, you can refer to it with ` x. This command does not change the current line. Default address: current line.

[line1 [,line2]] **m[ove]** *lineno*

Move the specified line(s) after *lineno*. The first moved line becomes the current line. Default address: current line.

n[ext]

Edit the next file in the argument list specified at startup. Default address: none.

n[ext]!

Move to the next file and overwrite the current buffer whether you've saved it or not. Default address: none.

n[ext] *file-list*

Replace the current list of files to edit with the specified *file-list* and edit the first file on new list. Default address: none.

n[ext]! *file-list*

Allow editing of new *file-list* even if you have not saved the current buffer. Default address: none.

n[ext] *+cmd file-list*

Execute *cmd* (which must not have spaces in it) after opening the first file in *file-list*. Default address: none.

Print the specified line(s) (*line1*, or the lines between *line1* and *line2*, or the next *n* lines) preceded by its line number. Including a *flag* after the command (either `p` for `print` or `l` for `list`) changes the display to the specified format. Default address: current line.

[line-spec] `o[pen]` *[/pat]* *[flags]*

Use `vi` commands at each addressed line. Specifying *pat* moves the cursor to the beginning of the string matched by the pattern. Including a *flag* after the command (either `p` for `print`, `l` for `list`, or `#` for `number`) changes the display to the specified format. Type `Q` to exit this mode. Default address: current line.

`pre[serve]`

Save the current editor buffer as though the system had just crashed. Use this command only in emergencies when a `write` command results in an error and you do not know how to save your work. Use the `recover` command or `ex -r` to recover after a `preserve`. Default address: none.

[line1[,line2]] `p[rint]` *[n]*

Print the specified line(s) (*line1*, or the lines between *line1* and *line2*, or the next *n* lines) displaying nonprinting characters as $^\wedge x$ and delete (octal 177) as $^\wedge ?$.

[line1[,line2]] `P[rint]` *[n]*

Print the specified line(s) (*line1*, or the lines between *line1* and *line2*, or the next *n* lines) displaying nonprinting characters as $^\wedge x$ and delete (octal 177) as $^\wedge ?$.

[lineno] `pu[t]` *[buffer]*

Put previously deleted or yanked lines after *lineno*. This moves lines with `delete` or copies lines with `yank`. Specifying *buffer* (a lowercase letter between `a` and `z`) retrieves text placed in that buffer with a `delete` or `yank` command. Default address: current line.

`q[uit]`

Leave `ex`. If you haven't saved all your changes, `ex` warns you and doesn't allow you to leave `ex`. `ex` also tells you if there are more files in the argument list. Normally, you should `write` your changes before doing a `quit`. Default address: none.

q[uit]!

Same as `quit`, but `ex` doesn't warn you if you haven't saved the current buffer. Default address: none.

[*lineno*] r[ead] [*file*]

Copy the text of *file* after *lineno* in the current buffer. If you don't supply *file*, it uses the current filename. If there is no current filename, *file* becomes the current name. It will not allow you to read in devices, but it will allow you to read in binary files.

If the file buffer is empty and there is no current file, `ex` treats this as an `edit` command. `0read` reads the file at the beginning of the buffer. It gives the same statistics as the `edit` command when it reads the file in. After a `read` command, the current line is the last line read (in `ex`) or the first line read (in open or visual mode). Default address: current line.

[*lineno*] r[ead] !*command*

Read the output of *command* into the buffer after *lineno*. There must be a blank or tab before the `!`. Default address: current line.

rec[over] *file*

Recover *file* after accidentally hanging up the phone, a system crash, or a `preserve` command. Default address: none.

rew[ind]

Start editing the files in the argument list, beginning with the first file you supplied when you started `ex`. Default address: none.

rew[ind]!

Same as `rewind`, but doesn't save the current buffer. Default address: none.

se[t]

The forms of this command are `set`. Print those options you've changed from their default settings.

`set all`	Print all the option values.
`set opt=val`	Give the value *val* (either a number or a string) to the option *opt*.
`set opt`	List the current setting of a string or numeric option.

`set` *opt*	Turn on an option that can be either off or on.
`set` `no`*opt*	Turn off an option that can be either off or on.
`set` *opt?*	List the current setting of an option that can be either off or on.

You can give more than one parameter to `set`; parameters are interpreted left-to-right. See "Command Option Summary and Descriptions" for the complete list. Default address: none.

`sh[ell]`

Create a new shell. Editing resumes when you cancel the new shell (using `exit`). Default address: none.

`so[urce]` *file*

Read *file* and run the (text-manipulation) commands in it. You can nest this command. Default address: none.

[line1 `[,` *line2]* `]` `s[ubstitute]` `/`*pat* `/`*repl* `[` `/`*suffix* `]`

s   Replaces the first instance of pattern *pat* with replacement pattern *repl* on each specified line. The suffixes are

g   (global.) Substitute *pat* with *repl* every time it appears on the specified lines. To make the substitution everywhere in the file, use the format `1,` `$s` `/`*pat/ repl/* `g`. You can also use `%` instead of `1,` `$`.

c   (confirm.) Print the line before making each substitution, marking the string to substitute with `^` characters. Type `y` to confirm the substitution, and type any other character if you don't want to make the substitution.

r   (replace.) Replace the previous replacement pattern from a substitution with the most recently mentioned regular expression; for example, from a search command.

You can split lines by substituting newline characters into them. You must escape the newline in *repl* by preceding it with a backslash (`\`). See "Regular Expressions and Searching" for other metacharacters available in *pat* and *repl*. Default address: current line.

[*line1* [*,line2*]] s[ubstitute] *suffix*

Omitting *pat* and *repl* repeats the last substitution. This is a synonym for the `&`
command, which is described later in this section. Using the `r` suffix (`sr`) replaces the
previous *pat* with the previous regular expression. This is a synonym for the `~`
command, which is described later in this section. Default address: current line.

[*line1* [*,line2*]] t *lineno flags*

`t` is a synonym for `copy`.

ta[g] *tag*

Start editing the file at *tag*, moving to another file if necessary. You must write the
current file, if you modified it, before giving a `tag` command. If you give the `tag`
command without specifying a *tag*, it uses the previous tag.

   Normally you use this command after using the `ctags`(1) command to create a tag
file. (See `ctags`(1) in *A/UX Command Reference.*) Default address: none.

una[bbreviate] *wd*

Delete *wd* from the list of abbreviations. When you type *wd*, it is not expanded.

u[ndo]

`undo` reverses the changes made by the last editing command, except `write`
or `edit`.

   `undo` marks the previous current line with `''`. If you restored a line, this becomes
the current line. If you didn't restore a line, the line before the last deleted line becomes
the current line.

unm[ap] *string*

Reverse the effect of a previous `map` command, removing the definition associated
with *string*. (Note that the `map` command only affects visual mode.) Default address:
none.

[*line1* [*,line2*]] v */pat/cmds*

A synonym for the variant form of a `global` command: runs *cmds* at each line *not*
matching *pat*. Default address: none.

ver[sion]

Give the current version of the editor and the last date the editor was changed. Default address: none.

[*line-spec*] vi [*type*] [*n*] [*flags*]

Enter visual mode at the specified line. The optional *type* argument (- ^ or .) specifies where the line is placed on the screen. If you omit *type*, the specified line is the first line on the screen. *n* specifies an initial window size; default is the value of the option window. Type Q to exit this mode. Default address: current line.

[*line1*[,*line2*]] w[rite] [*file*]

write writes changes back to *file*, printing the number of lines and characters written. Normally, you omit *file* and the text goes back where it came from. If you specify *file*, text is written to that file. By default, it writes the entire file.

   The editor writes to a file only if it is the current file, if it is creating the file, or if the file is actually a device (/dev/tty, /dev/null). Otherwise, you must give the variant form w! to force the write.

   If the file does not exist, ex creates it. This command does not change the current line. If there is an error while writing the current and edited file, the editor considers that there has been no write since the last change, even if the buffer had not previously been modified. Default address: current line.

[*line1*[,*line2*]] w[rite] >> *file*

Append buffer contents to *file*. Default address: current line.

w[rite]!*file*

Force a write to a file. This is helpful when you want to write to a file that already exists. Default address: none.

[*line1*[,*line2*]] w[rite] !*command*

Write the specified line(s) into *command*. Note that this is different from w! because a blank or tab must separate the w from the !. Default address: current line.

wq [*file*]

write followed by quit. Default address: none.

wq! [*file*]

The variant overrides checking on the `write` command, as `w!` does.

x[it] [*file*]

Write the buffer if there have been any changes, then quit the file. Default address: none.

[*line1* [,*line2*]] ya[nk] *buffer n*

`yank` places a copy of the specified line(s) in the named buffer. You can retrieve them with `put`. If you don't specify a buffer name, the lines go to a more volatile place (see the `put` command description). Default address: current address.

[*lineno+1*] z *n*

`z` prints the next *n* lines (default `window`).

[*lineno*] z*type n*

The `z` command determines where the current line appears on the screen. *type* is the character following the command and determines the positioning of the display on the screen. There are several different forms of this command:

[*lineno*] z or [*n*] z+	Print the next full screen of lines with the current (or specified) line at the top of the screen.
[*lineno*] z-	Print the screen with the current (or specified) line at the bottom.
[*lineno*] z.	Print the screen with the current (or specified) line at the center.
[*lineno*] z=	Print the screen with the current (or specified) line in the center, surrounded by lines of - characters.
[*lineno*] z^	Print the screen two windows before the current (or specified) line. Default address: current line.

!*command*

Send *command* to the shell to run. Within *command*, `%` and `#` are expanded as in filenames; `!` is replaced with the text of the previous command. Thus, `!!` repeats the last shell escape. If there is any such expansion, the expanded line is echoed. This command does not change the current line.

If you haven't written the buffer contents since the last change, `ex` prints a warning message before executing the command. A single `!` prints when the command completes. Default address: current line.

[line1[,line2]] !command

Supply the specified address (or address range) as standard input to *command*. The output then replaces the input line(s). Default address: current line.

[line-spec] =

Print the line number of the specified line. "Dot equals" (. =) gives the current line number. If no line is specified, the line number of the last line in the file is given. Does not change the current line. Default address: last line in file.

[line1[,line2]] < n flags* or *[line1[,line2]] > n flags

The less-than and/or greater-than signs (< and >) shift left or right a distance specified by the `shiftwidth` option. They shift only blanks and tabs and do not discard nonwhite characters in a left shift. The current line is the last line that changed due to the shifting. Default address: current line.

[line-spec] CONTROL-D

CONTROL-D scrolls through the file. You can specify the size of the scroll with the `scroll` option. The default is a half screen of text. Default address: current line.

[line1[,line2]] or RETURN

Print the addressed line(s). Pressing RETURN prints the next line. Default address: none.

[line-spec] & suffix n flags

The ampersand (&) repeats the previous substitute command on the current (or specified) line. If you set the `edcompatible` option, it retains the suffix; that is, if the previous substitute command was global, the ampersand repeats the substitution globally on the current line. Default address: current line.

[~ [line-spec]] ~suffix n flags

The tilde (~) replaces the previous replacement pattern from a substitution with the previous regular expression. Default address: current line.

[line1[,line2]] # [n] [flags]

Print the specified lines(s) (*line1*, or the lines between *line1* and *line2*, or the next *n* lines) preceded by its line number. Including a *flag* after the command (either p for print or l for list) changes the display to the specified format. Default address: current line.

5 Using the ed Line Editor
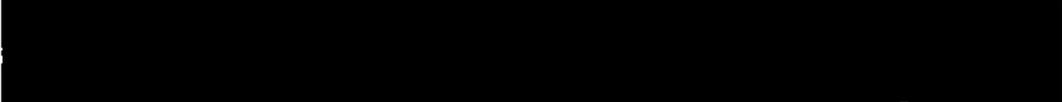
This chapter provides a detailed description of the commands and capabilities of the ed line editor.

What is `ed`?

The `ed` line editor is an interactive line-oriented text editor that uses your instructions to create and modify text files. A **line editor** moves through your file one line at a time and allows you to modify that line or to change another line or range of lines (indicated by line number). The `red` editor is a restricted version of `ed`. It is identical to `ed` except you can only edit files in the current directory and you cannot access shell commands.

The `ed` line editor is the only editor available in the A/UX Startup application (formerly known as `sash`). From the Startup application you can look at the A/UX file system while A/UX isn't running. With `ed` you can edit those files. This is very useful, for example, to edit the A/UX initialization file `inittab` before starting up A/UX. For more information on the Startup application, see *A/UX Local System Administration.*

◆ **Note** Except for the command you use to invoke the editor program, all commands discussed in this chapter are commands to `ed`. Do not confuse them with A/UX shell commands. ◆

Starting `ed`

There are two ways to start `ed`. The first and fastest way to start `ed` is from a CommandShell window. To start `ed` in this manner, type

`ed` *filename*

where *filename* may or may not already exist.
If a file by that name does not exist, you see the message

? *filename*

If a file by that name does exist, `ed` displays the character count on the screen.

You can also use the Commando command line interface to start the `ed` program and open a new or existing file. See Chapter 4, "Using Commando," in *A/UX Essentials* for more information about the Commando command line interface.

Displaying a prompt

You can use the P command to display a prompt on your screen. Type

P

The following appears on the left side of your screen:

*

   You type ed commands next to the asterisk (*) in the same way that you type shell commands next to the A/UX system prompt. To turn off the prompt, type the P command again.

Error messages

If ed doesn't understand something you type, it prints a question mark (?) on the screen.

For assistance in interpreting this error message, type

H

   The H (help) command explains the current error message and all subsequent ones. Typing the H command again turns off this feature.
   Alternatively, you can use the h command. This form of the help command explains only the current error message.

Inserting text

When you start ed, you open the editing buffer. The buffer corresponds to an empty file. It is a temporary work space, similar to a blank piece of paper. When you create a file, you must insert text into the buffer or read it in from another file and then save the new or modified data.

For example, when you give the command

`ed` *filename*

where *filename* is an existing file, `ed` makes a copy of this file and places it in the editing buffer. Any modifications or additions you make to this file are made on the copy, not on the original file.

The following example begins with inserting text in an empty buffer (editing existing files is discussed later).

To begin creating text, type

```
a
```

on a line by itself, and press RETURN. (The `a` command means append or add text lines to the buffer as they are typed in.)

Type the following text:

```
A journey of a
thousand miles
begins with a
single step.
.
```

As shown in the last line of this example, appending is stopped by typing a period (.) followed by RETURN. The period must be the first and only character on the line. This tells `ed` that you have finished adding text and are ready to give a command. Even experienced users sometimes forget to type the period when they have finished adding text. If `ed` seems to be ignoring your commands, type a period, and then press RETURN. You may find that some command lines in your text have to be removed.

After you finish appending, the buffer contains these four lines:

```
A journey of a
thousand miles
begins with a
single step.
```

To add more text, type

```
a
```

(RETURN), and continue typing.

Saving text

After you have added text to the buffer, you will want to save it. The `w` (write) command writes the contents of the buffer into a file. For example, if you type

```
w myfile
```

the buffer's contents are copied into a file named `myfile`.

If you named your file when you began your editing session, or if you are editing an existing file, you don't have to repeat that filename when you write the file. `ed` remembers the original filename you designated and automatically reuses it. For example, an editing session might look like this:

```
ed myfile
```
(*editing session*)
```
.
```
```
w
```

The file you edited is saved in a file named `myfile` when you type `w`.

You can also use the `w` command to save part of a file. The `w` command writes the lines you specify from the buffer to the permanent file. If no lines are specified, the `w` command writes the entire file. For example, the command

```
1,10w
```

saves the first ten lines of your file.

In another example, if you are editing your file `myfile`, and you give the command

```
1,10 w another.file
```

`ed` writes the first ten lines of your file `myfile` to the file `another.file`.

Note that when you assign a name to a file from within `ed` you must make sure that you do not have an existing file by that name. The write command replaces that file with the current buffer's contents *without* giving you a warning.

After writing the file, `ed` responds as follows:

```
57
```

This represents the number of characters (including blank spaces and end-of-line characters) that were written into the file.

◆ **Note** It's a good idea to write your text to a file every 10 or 15 minutes. ◆

Quitting `ed`

To quit `ed` after saving your text with the `w` (write) command, type

`q`

(followed by RETURN). For example, in the editing session described above, the following appears on the screen:

```
ed                      (start the editor program)
a                       (append )
A journey of a          (text)
thousand miles          (text)
begins with a           (text)
single step.            (text)
.                       (end append)
w text                  (write to a file named text)
57                      (character-count system response)
q                       (quit)
```

When you leave `ed`, the buffer is destroyed, and the system responds with its usual shell prompt character.

If you try to quit the editor without writing the buffer contents to a file, `ed` prints

`?`

on your screen.

If you don't want to save the changes to your file, typing `q` a second time (followed by RETURN) gets you out of `ed` and back to the shell without saving the changes you made since the last `w` command. If you want to save the changes to your file, type `w` and press RETURN.

Editing an existing file

After you have created and saved a file, you may want to edit it. Open the file using one of the methods described earlier, or type

`e` *filename*

When you use the `e` command to edit a file, `ed` replaces the contents of the buffer with the new file. If you were already working on a file in the buffer and you haven't written it yet, the `e` command destroys it without warning you.

If you forget the name of the file you have in the buffer, you can find out using the `f` (file) command. From within the editor, type

```
f
```

and the name of the file appears on the screen.

Displaying the contents of the buffer

To display all or part of the buffer on your screen, use the `p` (print) command. You must specify the line numbers where you want printing to begin and end. Separate these numbers with a comma in this format:

line1,line2`p`

Through this chapter, such line addressing is represented with the following:

line1,line2command

where *command* is `p` in this case. *line1,line2* indicates a range of addresses from *line1* to *line2*.

For example, to print the first ten lines of the buffer (lines 1 through 10), type:

```
1,10p
```

You can also tell `ed` to display the line numbers of the lines you specify with the `p` command. For example,

```
2,4pn
```

prints the following lines:

2   *text of line 2*

3   *text of line 3*

4   *text of line 4*

Suppose you want to print all the lines in the buffer. If you know the exact number of lines in the buffer, such as 30, you could type `1,30p`. However, if you don't know how many lines there are in your file, use the dollar sign (`$`). (The dollar sign refers to the last line of the file; see the section "Using Special Characters in `ed`.") To print all the lines in the buffer, type

```
1,$p
```

To stop printing, press the *interrupt* key (usually CONTROL-C). `ed` responds with

```
?
```

and waits for the next command.

To print the last line of the buffer, type

```
$p
```

You can print any single line by typing the line number. For example, typing

```
1
```

prints

```
A journey of a
```

which is the first line of the buffer.

In `ed`, the current line is the most recent line processed (in this case, the line last printed). If you type `p` again, `ed` prints line 1 again. The period character (or dot) always refers to the current line. It is a line number in the same way that `$` is. You can use dot in several ways—one possibility is to enter

```
.,$p
```

This prints everything from the current line to the last line of the buffer. In the example file `myfile`, these are lines 1 through 4.

Some commands move the current line to a new place in the file (that is, they change the value of dot); others do not. The `p` command resets dot to the number of the last line printed. For example

```
.,$p
```

sets dot to the last line in the buffer (line 4).

Dot is most useful in combinations such as

```
.+1     (this is equivalent to  .+1p)
```

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

```
.-1    (or  .-1p)
```

which means "print the line before the current line." This allows you to move backward through the buffer. Another useful example is

```
.-3,.-1p
```

which prints the previous three lines.

Don't forget that all of these commands change the value of dot. You can find out what dot is by typing

```
.=
```

This will print the line number of the current line. Pressing RETURN once prints the next line. It is equivalent to

```
.+1p
```

To summarize, you can precede p with zero, one, or two line numbers. If you don't specify a line number, p prints the current line (the line that dot refers to). If you specify one line number with or without the letter p, ed prints that line and makes it the current line. If you specify two line numbers separated by a comma and followed by p, ed prints everything from the first number to the last number, and sets dot to the last line printed. (The first number must be smaller than the second number—ed won't print backward.)

Typing the caret (^) or the minus sign (-) moves the current line back one line. These characters can be used in multiples; typing ^^^ or --- moves the current line back three lines. The minus (-) and caret (^) are the same as - 1p.

You can use line numbers with most ed commands, as you will see in the sections that follow.

Reading text into the buffer

If you want to add an existing file to the buffer without overwriting what is already there, use the `r` (read) command. The command

```
r new.file
```

adds the contents of the file `new.file` to the end of the file already in the buffer. If you type

```
e myfile
```

57   (system response)

```
r myfile
```

57   (system response)

the buffer now contains two copies of the same file:

```
A journey of a
thousand miles
begins with a
single step.
A journey of a
thousand miles
begins with a
single step.
```

Like the `w` and `e` commands, `r` prints the number of characters that it read into the buffer.

If you precede the `r` command with a line number or a dot (.), it reads a file and puts it after the specified place in the current buffer.

`.r` *filename*

reads the contents of *filename* into the buffer immediately after the current line. (In this context, dot is equal to the current line. This is different from the period character on a line by itself, which means that the text insertion is over.)

`3r` *filename*

reads the contents of *filename* into the buffer following line number 3.

The file in the buffer is not destroyed—it continues after the last line of the file you read in. For example, using the original file `myfile`

```
ed myfile
57                          (system response)
1                           (go to line 1)
A journey of a              (system response)
.r myfile
57                          (system response)
w
114                         (system response)
q
```

places this in your file:

```
A journey of a
A journey of a
thousand miles
begins with a
single step.
thousand miles
begins with a
single step.
```

Deleting text

The `d` (delete) command removes lines of text from the buffer. The `d` command uses the same format as the `p` command

line1,line2 d

For example, the command

```
4,$d
```

deletes everything from line 4 to the end of the buffer. In the preceding example, this deletion leaves us with three lines. We can check these lines by typing

```
1,$p
```

The last line, `$`, is now line 3. If you delete the last line (as in the preceding example), dot is set to `$`.

You can use the `d` (delete) command and the `p` (print) command together. For example, typing

```
dp
```

deletes the current line, prints the next line, and sets dot to the line printed.

Inserting text

The `i` (insert) command inserts one or more lines into the buffer. It is similar to the `a` command except that it places the text *before* rather than *after* the current or specified line—for example, typing

```
2i
```
one or more lines of text

.

inserts the text *before* the second line. If you don't specify a line number, the text is inserted before the current line. Dot is set to the last line inserted.

Experiment with the `i` and `a` commands to see how they operate. Verify that

line-spec`a`

text

.

appends *after* the given line, while

line-spec`i`

text

.

inserts *before* it, where *line-spec* indicates a single line number or a scanning command (such as a context search or regular expression) resulting in zero or more lines. If a line number isn't specified, the current line is assumed.

Changing text

The `c` (change) command changes the current line, replacing it with one or more lines. For example, to replace everything between the current line and the last line, type

```
.+1,$c
```
one or more lines of text
```
.
```

The text you type between the `c` command and the `.` command will overwrite the original text from the `.+1` line to the last line. This command is useful when you want to replace one line or several lines.

If you specify only one line, only that line is replaced. (You can type as many replacement lines as you like.) Notice that you end your changes by typing a period (`.`) at the beginning of a line—this is the same way you stopped adding text with the `a` command.

The `c` command can also be thought of as a combination of the `d` command followed by the `i` command. Experiment to verify that

line1,line2 `d`

```
i
```
text
```
.
```

is the same as

line1,line2 `c`

text
```
.
```

If you don't specify a line number, `c` replaces the current line. When you finish making changes, dot is set to the last line you inserted.

Substituting text

One of the most important `ed` commands is the `s` (substitute) command.

This command changes words or characters and can be used to correct spelling mistakes and typing errors.

Suppose that line 1 is

```
A journy of a
```

You can change `journy` to `journey` by typing

```
1s/ny/ney/
```

This says: in line 1, change `ny` to `ney`. Since `ed` doesn't print the change automatically, type

```
p
```

to make sure the substitution worked. You should see

```
A journey of a
```

When you include the `p` command on the same line as the substitute command

```
s/journy/journey/p
```

`ed` prints the line that just changed.

The general format of the substitute command is

line1, *line2* `s`/ *change this*/ *to this*/

The characters between the first and second slashes (*change this*) are replaced by the characters between the second and third slashes (*to thi*s). This substitution takes place on *all* lines between *line1* and *line2*. However, only the *first* occurrence on *each* line is changed. To change *every* occurrence, on *each* line, add `g` (global) (see "Global Commands") to the `s` command, like this:

line1, *line2* `s`/ *buckwheat*/ *farina*/`g`

The rules for line numbers are the same as those you learned for the `p` (print) command. However, if the `s` command can't find the characters you asked it to change, the cursor stays in the current position. The `ed` line editor alerts you when this happens by printing `?` on the screen.

As an example of a substitution, you could type

```
1,$s/speling/spelling/
```

to correct the *first* instance of `speling` on each line. (This is useful for people who make the same mistake consistently.)

If you don't specify a line number, `s` assumes you want to make the substitution on the current line. For example, you could type

`s`/ *buckwheat*/ *farina*/`p`

This corrects a mistake on the current line and then prints it to verify that the substitution worked.

You may have noticed that the `s` command resets the current line. You can also type

s/ *buckwheat*//

This replaces *buckwheat* with nothing—in other words, it removes *buckwheat*. This is useful for deleting extra words in a line or removing extra letters from words.

For example,

```
Thisxx is an example of substitution
```

can be corrected by typing

```
s/xx//
```

The line now reads

```
This is an example of substitution
```

The `//` (two adjacent slashes) mean "no characters," *not* a blank.

Experiment with the `s` command. For example, type

```
a
the other side of midnight
.
s/the/meet me on the/p
```

This produces the following:

```
meet me on the other side of midnight
```

Remember that the `s` command changes only the first occurrence. You can change all occurrences on a line by adding `g`.

Try using characters (except blanks and tabs) other than slashes to set off the two sets of characters in the `s` command. For example, try typing

```
s'the'meet me on the'p
```

This works exactly the same as using a slash.

However, strange results are produced by using the backslash (\ ) character. See "Using Special Characters in `ed`," later in this chapter, for more information.

Global commands

The `g` (global) command performs an operation on all lines that match a specified string or regular expression. See Chapter 4, "Using the `ex` Line Editor," for information on regular expressions; in this chapter we use the word *string* to mean a string of characters or a regular expression. For example,

```
g/speling/p
```

prints all lines that contain `speling`. The command

```
g/speling/s//spelling/gp
```

replaces `speling` with `spelling` each time it occurs (even if it occurs more than once in a line), then prints each corrected line.

Compare this to

```
1,$s/speling/spelling/gp
```

This prints only the last line substituted.

You can use several commands at a time with `g`. Just remember to end every line but the last with a backslash (`\`). For example,

```
g/xxx/-1s/abc/def/\
.+2s/ghi/jkl/\
.-2,.p
```

makes changes in the lines before and after each line containing `xxx`, then prints all three lines.

The `G` (interactive global) command finds a line that matches a specified string, prints the line, and waits to accept a command. After executing the command, it searches for the next line that matches the specified string, and so on. For example,

```
G/speling/
```

prints the first line that contains the string `speling`. If you wish to change the string at that point, you can enter the command

```
s/speling/spelling/p
```

which replaces `speling` with `spelling` and prints the corrected line. After printing the corrected line, `ed` searches for the next instance of `speling`. If found, it prints the line that contains this string, and waits for you to enter a command. The

command you enter does not have to be the same command you entered last time; for example, if `ed` finds another instance of `speling`, you could enter the command

```
s/speling/misspelling/p
```

or any single `ed` command other than the `a`, `c`, `i`, `g`, `G`, `v`, or `V` commands.

The `v` command is the same as `g` except that it runs the commands on lines that don't match the string or regular expression. For example,

```
v/ /d
```

deletes every line that does not contain a blank. Similarly, the `V` command is the same as `G`, but finds and prints lines that don't match the specified string or regular expression.

Searching for a character string

When you master the substitute command, you may want to try another important feature of `ed`—context searching. Context searching looks for a string of characters and, when it finds it, makes that line the current line.

Suppose you have these three lines in your buffer:

```
Little Miss Muffet
sat on a tuffet
eating her kurds and whey.
```

If you want to locate the misspelled word *kurds,* you could type `3`. However, if the buffer contained several hundred lines and you had been deleting and rearranging lines, you might have a difficult time locating the line you wanted. Context searching lets you find a line by specifying some context (unique text) in it.

To search for a line that contains a particular string of characters, type

/ *string of characters/*

For example,

```
/kurds/
```

locates the next occurrence of *kurds*. It also makes that line the current line and prints it for verification.

Next occurrence means `ed` starts looking for the string at the line following the current line (`.+1`) and searches to the end of the buffer. Then it searches from line 1 to the line it started searching at (dot). That is, the search wraps around from `$` to `1`. It scans all the lines in the buffer until it either finds the desired string or gets back to dot again. If `ed` can't find the characters, it types the error message

```
?
```

Otherwise, it prints the line it found.

You can search for the desired line and make a substitution to it in the same command, like this:

```
/eating/s/kurds/curds/p
```

This tells `ed` to search for the line with the word `eating`, substitute `curds` for `kurds`, and then print the new line. When it has finished, `ed` prints this:

```
eating her curds and whey.
```

You can repeat a context search. For example,

/string/

finds the next occurrence of *string*. If this is not the line you want, you can search for the next occurrence by typing

```
// or /
```

This stands for the previous context search expression and differs from the use of `//` as a null argument in the `s` command.

This abbreviation can also be used as the first string of the `s` command. For example,

```
/string1/s//string2/
```

finds the next occurrence of `string1` and replaces it with `string2`. Similarly,

```
?? or ?
```

scans backward for the previous expression.

You can use context searches instead of line numbers to find a desired line or to specify a range of lines to be affected by some other command, such as `s`.

For example, suppose the buffer contains these four familiar lines:

```
A journey of a
thousand miles
begins with a
single step.
```

The following context search expressions all refer to the same line (line 2):

```
/journey/+1
/thousand/
/step/-2
```

To make a change in line 2, you can type

```
/journey/+1s/thousand/hundred/
```

or

```
/thousand/s/thousand/hundred/
```

or

```
/step/-2s/thousand/hundred/
```

You could print all four lines by typing either

```
/journey/,/single/p
```

or

```
/journey/,/journey/+3p
```

The first of these might be better if you don't know how many lines there are. A context search expression is the same as a line number, so it can be used wherever you would use a line number.

Moving text

The `m` (move) command moves lines from one place to another. For example, to move the first four lines of the buffer to the end, type

`1,4m$`

The general case is

line1, *line2* `m` *lineno*

The text is moved after the specified line number (*lineno*). You can use context searches instead of line numbers. For example, if you have the following text in your buffer,

First paragraph

. . .

end of first paragraph
Second paragraph

. . .

end of second paragraph.

you could reverse the two paragraphs by typing

/ *Second* / , / *end of second* / `m` / *First* / `-1`

The `-1` was used because the text is moved *after* the line specified. Dot is set to the last line moved.

Using special characters in `ed`

You may have noticed that some characters (such as `.`, `*`, `$`) change the meaning of context searches and the `s` command. This is because these characters have special meanings for `ed`.

The following is a complete list of these special characters:

`^ $ * [ ] & \`

These are described in the sections that follow.

The period (.) character

In a context search or the first string of the substitute command, the period character (.)
signifies *any* character.

Although this is the same character as "dot," its meaning is different in this context.
To avoid confusion, we call it *dot* when it means "current line" or "line most recently
changed" and *period* when it means "any character."

```
/x.y/
```
means
x any-character `y/`

This command will find all instances of `x` followed by any character followed by `y`,
including the following:

```
x+y
x-y
x y
x.y
xAy
```

The caret (∧) character

The caret character (^) signifies the beginning of a line. For example,

```
/^bunny/
```

finds *bunny* only if it is at the beginning of a line. That is, it will find

```
bunny
```

but not

```
bugs bunny
```

The dollar sign ($) character

The dollar sign character (`$`) is the opposite of the caret; it means the end of a line.

The expression

```
/bunny$/
```

finds *bunny* only at the end of a line.

```
/bunny$/
```

finds a line containing only bunny and

```
/^.$/
```

finds a line containing one character.

The asterisk (*) character

The asterisk character (`*`) is the repetition character. For example, `a*` means "zero or more a's." `.*` means "any character repeated zero or more times."

For example,

```
s/.*/stuff/
```

changes an entire line to `stuff`, and

```
s/.*,//
```

deletes all the characters in the specified line up to, and including, the last comma. Note that `*` finds the longest possible match, so this example matches the last comma rather than the first.

The bracket ([ ]) characters

The left bracket character ([ ) is used with the right bracket character ( ]) to enclose "character classes." For example,

```
/[0123456789]/
```

searches for any single digit. This can be abbreviated as

```
[0-9]
```

Brackets can also be used to contain a character class that represents the alphabet; for example,

```
[A-Z]
```

searches for any uppercase character, and

```
[a-z]
```

searches for any lowercase character.

The ampersand (&) character

The ampersand character (&) means "whatever was matched on the left-hand side." (The ampersand only has this meaning on the right-hand part of a substitute command.)

Suppose the current line contains

```
Drop the gun
```

and you want to put parentheses around it. You can accomplish this using the command

```
s/.*/(&)/
```

This tells ed to match the whole line (. *) and replace it by itself (&) surrounded by parentheses so that it appears as follows

```
(Drop the gun)
```

The ampersand can be used several times in a line. Using the preceding sample text,

```
s/.*/&? &!!/
```

produces

```
Drop the gun? Drop the gun!!
```

You don't have to match the whole line. For example, if the buffer contains

```
it's starting to hit me
```

you could type

```
/me/s//& like a two ton heavy thing/
```

to produce

```
it's starting to hit me like a two ton heavy thing
```

The sequence `/me/` found the desired line; the sequence `//` found the same word in the line; and the `&` saved you from typing `me &` again.

The `&` is a special character only in the replacement text of a substitute command.

The backslash (\) character

If you have to use one of the special characters listed above without its special meaning in a substitute command, precede it with a backslash (`\`). For example,

```
s/\.T//
```

replaces the first occurrence of a `.T` with nothing (`//`) on the current line (in other words, it deletes it). If the period (`.`) were not preceded by the `\`, the result would have been that the first instance of `H` preceded by *any* other character would have been deleted on the current line.

Command summary

In the following summary, *line-spec* indicates a single line number or a search command (such as a context search or regular expression) resulting in zero or more lines; *line1, line2* indicates a range of addresses from *line1* to *line2*. If you don't specify an address, the current line is the default (unless otherwise noted). Portions of a command enclosed in brackets ([ ]) are optional.

[*line-spec*] a RETURN **[*text*]** RETURN.

Append text after the current line or after the line number specified. To stop adding text, type a period (.) at the beginning of a line, and press RETURN. Dot is set to the last line appended.

[*line-spec*] c RETURN **[*text*]** RETURN.

Change the specified lines to the new text that follows. To stop replacing text, type a period (.) at the beginning of a line, and press . RETURN. If you don't specify a line, the current line is replaced. Dot is set to the last line changed.

[*line1, line2*] d

Delete the specified lines. If you don't specify a line, the current line is deleted. Dot is set to the line after the last deleted line. If you delete the last line in the buffer, dot is set to the new last line.

e *file*

Edit a new file from within ed. The previous contents of the buffer are destroyed, so save your work before you edit a new file with e.

f [*file*]

Print the current filename. This is the file ed assumes you mean if you don't specify a file. To change the current filename, type f file.

[*line1, line2*] g /*string* /*command*

Execute commands globally, on the entire file (by default). The . g / *x* / . command runs on lines containing the string *x*.

[*line1* , *line2*] G / *string* [/]

Interactive global command. ed first marks every line that matches the given regular expression or string. Then, for every such line, that line is printed, dot is changed to that line, and any *one* command (other than one of the a, c, i, g, G, v, and V commands) may be input and is run. After the execution of that command, the next marked line is printed, and so on; a RETURN acts as a null command (no action is performed); an & causes the reexecution of the most recent command runs within the current invocation of G. Note that the commands input as part of the execution of the G command may address and affect *any* lines in the buffer. The G command can be canceled by an *interrupt*. A command that causes an error cancels the G command.

h

The h (help) command gives a short error message that explains the reason for the most recent ?.

H

The H (Help) command prints error messages for all subsequent ? diagnostics. This command toggles error message printing on and off.

[*line-spec*] i RETURN [*text*] RETURN .

Insert text before the specified line or the current line. To stop inserting text, type a period (.) at the beginning of a line, and press RETURN. Dot is set to the last line inserted.

[*line1* , *line2*] j

Join contiguous lines by removing appropriate newline characters.

[*line-spec*] k*x*

Mark addressed line with name *x*, which must be a lowercase letter. The address *x* then refers to this line; dot is unchanged.

[*line1* , *line2*] m*lineno*

Move the text originating between *line1* and *line2* to follow *lineno*. Dot is set to the last line moved.

[*line1* , *line2*] n

For the current line or for each line in the range specified by "*line1* , *line2*," print the line number, followed by a tab, followed by the text of the line(s).

[*line1* , *line2*] p

Print the specified lines. If you don't specify any line number, p prints the current line. Pressing RETURN prints the next line.

P

Turns prompting on and off. The P command alternately turns this mode on and off; initially it is off.

q

Quit ed. No automatic write of a file is done. If changes have been made in the buffer since the last w command, ed responds with ?. If you don't want to save your changes, type q RETURN again.

Q

Quit without checking to see if changes have been made in the buffer since the last w command.

[*line-spec*] r *file*

Read a copy of *file* in at the specified location. If no line number is specified, it reads the file in at the end of the buffer. Dot is set to the last line read.

[*line1* , *line2*] s / *string1* / *string2* [/]

Substitute one string for another string at a specified location. 1, $s / *string1* / *string2* / g substitutes *string2* for every instance of *string1* in the file. The s command changes only the first occurrence of *string1* on a line. To change all occurrences, type g at the end of the command. Dot is set to the last line in which a substitution took place; if no substitution took place, dot is not changed.

[*line1* , *line2*] t*lineno*

Put a copy of the addressed lines after address *lineno* (which may be 0); dot is left at the last line copied.

u

Undo last command; nullifies the effect of the most recent command that modified anything in the buffer.

line1, line2] v/ string / command

Execute *command* only on lines *not* containing *string*. By default the v command operates on the entire file.

line1, line2] V/ string [/]

Interactive global command marks each line not containing *string* and then allows you to perform commands on each of these lines. By default the v command operates on the entire file.

line1, line2] w file

Write the buffer into the specified file. Dot is not changed. By default the w command writes the entire file.

x

Request an encrypted key string from the standard input. Subsequent e, r, and w commands encrypt and decrypt the text with this key by the algorithm of crypt(1). An explicitly empty key turns off the encrypt function.

[.] =

"Dot equals" prints the current line number. = by itself prints the line number of the last line in the file.

! shell-command

Temporarily escape to the A/UX shell to run the specified command. ! *shell-command* runs *shell-command* in the shell and then returns you to the editor.

/ string [/]

Search through the file for *string* and print the line containing it. The search starts at the line after the current line, reads to the end of the buffer, then wraps around to line 1 and searches to the original line. If *string* is located, dot is set to the line where the string is found.

? string [?]

Search backward through the file for *string* and print the line containing it. The search begins at the line before the current line, reads backward to the start of the file, then wraps around to the end of the file and searches backward to the original line. If *string* is located, dot is set to the line where the string is found.

6 Using the `sed` Stream Editor

This chapter provides a detailed description of the commands and capabilities of the `sed` stream editor.

What is `sed`?

The `sed` stream editor is useful for creating filters for batch editing. Batch editing means that you run a file through a series of predetermined editing commands (filters) that automatically edit the file with no supervision required. You can also use `sed` for:

- Editing large files that cannot be contained in a buffer. The size of a file to be edited with `sed` is limited only by the amount of secondary storage. Only a few lines of the current input file are in physical (volatile) memory at one time, and no temporary files (buffers) are used.

- Performing complicated editing sequences on any size file. The `sed` editor is most commonly used in shell scripts, where complicated editing requests can be stored, edited, and applied to the input file(s) as a command.

- Efficiently performing multiple global editing commands in one pass. The `sed` program running from a command file is faster and more efficient than an interactive editor like `ex`, even when `ex` is also running from a command file.

- Performing transformations on a data stream as part of a pipe or a shell script.

Note that `sed` does not recognize certain commands provided by an interactive editor. For example, `sed` does not provide relative addressing. Because it operates on one line at a time, `sed` cannot move backward or forward relative to the current line in a file. In addition, `sed` does not inform you about the effects of your commands, or allow you to undo them immediately.

Overall operation

By default, `sed` copies standard input to standard output, performing zero or more editing actions on each line before writing it to the output. Editing actions are specified by `sed` editing commands, described in the next section. You specify the lines to be affected by these commands by addresses, either context addresses or line numbers.

You never modify an input file directly; instead, changes are written to the standard output. If this output is redirected to a file, then the new file contains the modifications created by your editing actions. Then you may, if you wish, replace the original file with this new file.

Command options

The following sections describe the function and syntax of the command options for the sed stream editor.

Command syntax

The command syntax for the sed editor is

```
sed [-n] -e 'command-line-script' [-e 'command-line-script']...
      [-f sfile]... [file...]
```

or

```
sed [-n] -f sfile... [-f sfile]... [-e 'command-line-script']... [file...]
sed [-n] 'command-line-script' [file...]
```

◆ **Note** sed must be invoked with at least one -e or -f option; however, if only "-e 'command-line-script' " is used, the -e may be omitted. ◆

sed can be invoked in any of the following ways:

```
sed 'command-line-script' file

sed -e 'command-line-script' file

sed -n -e 'command-line-script' file

sed -f sfile file

ed -n -f sfile file

sed -e 'command-line-script' -f sfile file
```

Table 6-1 provides a summary of all `sed` command options.

Table 6-1 `sed` command options

Option	Description
`-n`	(no-copy.) Copy only those lines explicitly specified either by `p` (print), `i` (insert), or `a` (append) commands or `p` arguments after `s` (substitute) commands.
`-e`	(expression.) The *command-line-script* argument is an "expression" (inline `sed` command(s) using the syntax of regular expressions and enclosed in single or double quotes) to be run on the input stream. There may be more than one `-e` (with its corresponding expression) on a command line. If the newline characters are preceded by an ESCAPE character, there may be more than one line in an expression. The `-e` itself may be omitted if there is only one expression and no `-f` option is present.
`-f`	(source file.) The *sfile* argument is a file containing `sed` commands, one to a line. There may be more than one `-f` option specified on the command line. If multiple `-f` command file arguments are given, the commands they contain are run in the order specified.

The input *files* may be omitted; in that case, `sed` takes its input from the standard input. Note that `sed` does not accept the "`-`" construct used in other programs (for example, `awk`) to indicate the standard input. If you must apply a sequence of `sed` commands to some files and then to the standard input, you can use the following command:

`cat` *files* `- | sed -f` *sfile*

Using commands

Editing commands are specified on the `sed` command line. They can either be embedded inline (with the `-e` option) or enclosed in a file and provided as an argument to the `-f` option. The following are examples of `sed` usage:

```
sort chap.1 |

sed -e 's/\.dc\./.dec./' -e 's/\.3b\./.u3b./'
```

This sorts the contents of `chap.1` and performs substitutions on the first instance of `.dc.` and `.3b.` in each line; the results are written to standard output.

Note that

```
sed -e 's/\.dc\./.dec./' -e 's/\.3b\./.u3b./'
```

is equivalent to

```
sed '
    s/\.dc\./.dec./
    s/\.3b\./.u3b./
'
```

In this chapter, we use the first form, which employs the `-e` option. These may be replaced with the second form if you prefer.

With `-e`, you may also separate editing commands with a semicolon. For example,

```
sed -e 's/ \.dc\./.dec./;s/ \.3b\./.u3b./'
```

is equivalent to the above examples.

The command form

```
sed -e 's/ \.dc\./.dec./
s/ \.3b\./.u3b./' chap.1
```

performs the same substitutions as the preceding command on a file named `chap.1`; the results are written to standard output.

◆ **Note** When using `sed` in the C shell, newline characters must be preceded by an *escape* character (backslash) even when enclosed in single quotes. ◆

The command

```
sed -e 's/,/ /g' chap.1 > temp
```

replaces every comma (,) in `chap.1` with a space; the modified file is contained in `temp`.

If you put the following `sed` commands into a file named `cmd.file`

```
s/ \.dc/.dec./g
s/ \.3b\./.u3b./
s/,/ /g
```

then you can use the following command:

```
sed -f cmd.file chap.1 > temp
```

This performs substitutions on `chap.1`; the modified file is contained in `temp`.

You can also use the command

```
sed -n -f cmd.file chap.1
```

to perform substitutions on `chap.1` and write the modified `chap.1` to standard output.

Before any input file is opened, all editing commands are compiled in the order encountered (also the order in which they are attempted at execution time) into a form that will be moderately efficient during the execution phase. In the execution phase the commands are actually applied to lines of the input file.

Editing command syntax

The general editing command syntax is

[*line-spec*] *command* [*arguments*]

The *line-spec* (line specification) and the *arguments* are optional, although either of these may be required according to the command given. *line-specs* may be line number(s) or context addresses in the form

[*line1*[, *line2*]]

or

[/*pattern*[/][, /*pattern*[/]]

In the first case, if one line number is specified, `sed` performs the editing command on that line; if two line numbers are specified, `sed` performs the editing command on the range of lines between *line1* and *line2*, inclusive. In the second case, if one context address is specified, `sed` performs the editing command only on lines containing that *pattern*; if two context addresses are specified, `sed` performs the editing command on

all lines between the first *pattern* and the second *pattern*, inclusive. After it recognizes the second pattern, `sed` searches for the first pattern again. If found, it begins performing the editing command again until it recognizes the second pattern, and so on. Any number of blanks or tabs may separate *line-specs* from the command; blanks and tab characters at the beginning of lines are ignored.

Command application order

Commands are applied one at a time, in the order encountered; the input to each command is the output of all previous commands.

This default linear ordering can be changed by the `t` (test substitution) and `b` (branch) control-flow commands. When the order of application is changed by these commands, it is still true that the input line to any command is the output of any previously applied commands.

Pattern space

The **pattern space** is the buffer the `sed` commands operate on. Ordinarily, pattern space is one line of the input text, but more than one line can be read into the pattern space by using the `N` command or the `G` command.

Addressing

Input file lines to be affected by your editing commands are specified by *line-specs*. These *line-specs* can be either line numbers or context addresses. If no *line-spec* is present, the command is applied to every line in the input file.

Multiple commands can be applied to a single *line-spec* by grouping commands with braces in the following format:

line-spec   {

                command-list

        }

Line number addresses

A line number is a positive decimal integer that is measured in increments (by an internal counter) as each line is read from the input. A **line number address** corresponds to the value of the internal line counter. As a special case, the $ character matches the last line of the last input file.

◆ **Note** The line counter runs cumulatively through multiple input files. It is not reset when a new consecutive input file is opened. ◆

Commands can be preceded by zero, one, or two addresses. It is an error when a command has more addresses than allowed.

If a command has zero addresses, it is applied to every line in the input.

If a command has one address, it is applied to all lines that match that address.

If a command has two addresses separated by a comma, it is applied to the first line that matches the first address and to all subsequent lines up to, and including, the first line that matches the second address. An attempt is made on subsequent lines to match the first address again, and the process is repeated.

Context addresses

A **context address** is a regular expression enclosed by matching delimiters. Any character may be selected as the expression delimiter (for example, / *pattern*/ or % *pattern*%). sed recognizes regular expressions that have the following construction:

- An ordinary character is a regular expression and matches that character.
- A caret (^) at the beginning of a regular expression matches the beginning of a line.
- A dollar sign ($) at the end of a regular expression matches the end of a line.
- The (\n) character matches an embedded newline character in the pattern space but not the newline character at the end of the pattern space. Newline characters may be embedded by using the N command or the G command.
- A period (.) matches any character except the terminal newline character of the pattern space.

- A regular expression followed by an asterisk (`*`) matches any number (including zero) of adjacent occurrences of the regular expression it follows.

- A string of characters in square brackets (`[]`) matches any character in the string and no others. For example, `[abc]` matches the single-character strings `a`, `b`, and `c`. The characters may also be specified as a range using the format

  `[a-z]`

  which will match any lowercase character. If, however, the first character of the string is a caret (`^`), the regular expression matches any character except the characters in the string and the terminal newline character of the pattern space. The caret is the only metacharacter recognized within the square brackets. If (`]`) needs to be in the string enclosed in square brackets, it should be the first non-metacharacter. Thus, for example,

  `[]...]`   includes `]`

  `[^]...]`   does not include `]`

  In both cases, a range may be specified by using a hyphen (for example, `[A-Z]` or `[0-9]`).

- A concatenation of regular expressions is a regular expression that matches the concatenation of strings matched by the components of the regular expression.

- A regular expression between the sequences `\(` and `\)` is identical in effect to the unadorned regular expression, but has side effects, which are described under the substitute command (`s`) later in this section.

- The expression `\d` (where *d* is a digit, 0 through 9) refers to the string of characters found earlier in the same pattern by an expression using the `\(` and `\)` construction. The `\(` and `\)` sequences act just like those in the other A/UX editors, and are used to establish "fields" or sections in a line of text (or all lines of text) in a file. For example, suppose a file contained the following list of names:

```
Dick Powell
William Powell
Hosken Powell
Jane Powell
```

  The following expression reverses the order of the names, while placing a comma and a space between each first name and last name:

```
s/\([A-Z].* \)\([A-Z].*\)/\2, \1/
```

This command writes a new list:

```
Powell, Dick
Powell, William
Powell, Hosken
Powell, Jane
```

For another example, the following expression matches a line beginning with two repeated occurrences of the same string separated by a space:

```
/^\(.*\) \1/
```

- A null regular expression standing alone (for example, `//`) is equivalent to the previous regular expression.
- Special characters `^`, `$`, `*`, `\`, and `/`, when used as literal characters, must be preceded by a backslash (`\`).
- For a context address to match, the whole pattern within the input address must match some portion of the pattern space.

Examples

Let us consider more examples of using `sed`. First, create a text file named `poem` that contains the following lines:

```
In Xanadu did Kubla Khan
A stately pleasure dome decree:
Where Alph, the sacred river, ran
Through caverns measureless to man
Down to a sunless sea.
```

Examples in this chapter use this text except where noted. The following example shows the output of a `sed` command using line number addressing. The command

```
sed -e '2q' poem
```

copies the first two lines of the input and quits. The output on your screen will be

```
In Xanadu did Kubla Khan
A stately pleasure dome decree:
```

On the same input text, the following lists the matches resulting from several `sed` commands using context addressing:

`/an/`	matches lines 1, 3, and 4
`/an.*an/`	matches line 1
`/^an/`	matches no lines
`/./`	matches all lines
`/\./`	matches line 5
`/r*an/`	matches lines 3 and 4
`/\(an\).*\1/`	matches line 1

Command summary

In the following summary, *line-spec* indicates a single line number or a context address. *line1*, *line2* indicates a range of addresses from *line1* to *line2*. If you don't specify an address, the commands are applied to all lines in the file (unless otherwise noted).

Line-oriented commands

Table 6-2 summarizes all of `sed`'s line-oriented commands. The commands in this section apply to the entire line (or lines) currently stored in the pattern space.

Table 6-2 `sed` line-oriented commands

Command	Description
[*line1* [**,** *line2*]] d	(delete.) The `d` command deletes from the file (does not write to the output) those lines matched by its addresses. It also has the side effect that no further commands are attempted on the remains of a deleted line. As soon as the `d` command is run, a new line is read from the input, and the list of editing commands is restarted from the beginning on the new line.
[*line1* [**,** *line2*]] n	(next.) The `n` command reads the next line from the input, replacing the current line. The current line is written to standard output. The list of editing commands continues following the `n` command.
[*line-spec*] a \ Return *text*	(append.) The `a` command causes the text argument to be written to the output after the line matched by its address.The `a` command is inherently works on multiple lines; `a` must appear at the end of a line, and *text* may contain any number of lines. To preserve the one-command-to-a-line fiction, interior newline characters must be hidden by a backslash character (\ ) immediately preceding the newline character.
	The *text* is deleted by the first newline character not immediately preceded by a backslash. Once an `a` command is successfully run, text will be written to the output regardless of what later commands do to the line that triggered it. Even if that line is deleted, text will still be written to the output. The *text* is not scanned for address matches, and no editing commands are attempted on it. The `a` command does not cause a change in the line number counter.

Table 6-2 `sed` line-oriented commands (*continued*)

Command	Description
[*line-spec*] i \ Return text	(insert.) The `i` command causes the text argument to be written to the output before the line matched by its address. The `i` command inherently works on multiple lines; `i` must appear at the end of a line, and text may contain any number of lines. To preserve the one-command-to-a-line fiction, interior newline characters must be hidden by a backslash character (\ ) immediately preceding the newline character. The text is deleted by the first newline character not immediately preceded by a backslash. Once an `i` command is successfully run, text will be written to the output regardless of what later commands do to the line that triggered it. Even if that line is deleted, text will still be written to the output. The text is not scanned for address matches, and no editing commands are attempted on it. The `i` command does not cause a change in the line number counter.
[*line1*[, *line2*]] c \ Return *text*	(change.) The `c` command deletes lines selected by its addresses and replaces them with the lines in the text argument. Like `a` and `i`, `c` must be followed by a newline character hidden by a backslash; interior newline characters in text must be hidden by backslashes. The `c` command may have two addresses and therefore select a range of lines. If it does, all lines in the range are deleted, but only one copy of text is written to the output, not one copy per line deleted.
	As with `a` and `i`, text is not scanned for address matches, and no editing commands are attempted on it. It does not change the line number counter. After a line has been deleted by a `c` command, no further commands are attempted on the corpse. If text is appended after a line by `a` or `r` commands and the line is subsequently changed, the text inserted by the `c` command will be placed before the text of the `a` or `r` commands. (The `r` command is described later.)
	Leading blanks and tabs disappear from text inserted in the output by the `a`, `i`, and `c` commands. To get leading blanks and tabs into the output, precede the first desired blank or tab with a backslash. The backslash will not appear in the output.

The following example shows line-oriented `sed` commands used on the standard input file poem.

If the file script contains the lines

```
n
a\
  XXXX
d
```

the command

`sed -f script poem > ` *output.file*

produces an *output.file* that contains the following lines:

```
In Xanadu did Kubla Khan
XXXX
Where Alph, the sacred river, ran
XXXX
Down to a sunless sea.
```

The substitute command

The substitute command uses the following syntax:

[*line1*[, *line2*]]s *pattern replacement flags*

The `s` command replaces the part of a line selected by *pattern* with *replacement*. It can be read "substitute for pattern, replacement." The command arguments are described as follows:

pattern   The *pattern* argument is a regular expression, like the patterns in context addresses. The only difference between *pattern* and a context address is that the context address must be delimited by slash (`/`) characters; *pattern* may be delimited by any character other than space or newline. By default, only the first string matched by *pattern* is replaced unless the `g` flag (below) is invoked.

replacement The *replacement* argument begins immediately after the second delimiting character of *pattern* and must be followed immediately by another instance of the delimiting character. Thus, there are exactly

three instances of the delimiting character. The *replacement* is not a pattern, and the characters that are special in patterns do not have special meaning in *replacement*. Instead, the following other characters are special:

&   is replaced by the string matched by *pattern*.

d   is replaced by the substring *d* (*d* is a single digit), matched by parts of *pattern*, and enclosed in \\(and \\). If more than one substring occurs in *pattern*, the substring *d* is determined by counting opening delimiters (\\(). As in *pattern*, special characters may be made literal characters by preceding them with a backslash (\\).

flags   The *flags* argument may contain the following:

g   (global.) Substitute *replacement* for all instances of *pattern* that do not overlap in the line. After a successful substitution, the scan for the next instance of *pattern* begins just after the end of the inserted characters. Characters put into the line from *replacement* are not rechecked.

p   (print.) Print the line if a successful replacement was done. The p flag causes the line to be written to the output if a substitution was actually made by the s command. If several s commands, each followed by a p flag, successfully substitute in the same input line, multiple copies of the line will be written to the output, one for each successful substitution. Note that unless the -n flag option is used, each line will be echoed automatically to standard output. In addition, each line affected by the p flag will be echoed as well, causing multiple copies to be written to standard output.

w *file* (write to file.) Write the line to a file if a successful replacement was done. A single space must separate w and file. The w flag causes lines that are actually substituted by the s command to be written to a file named by file. If file exists before sed is run, it is overwritten; if not, it is created. The possibilities of multiple, somewhat different copies of one input line being written are the same as for p. A maximum of ten different filenames may be mentioned after w flags and w commands.

The command

```
cat poem | sed -e 's/to/by/w changes'
```

produces an output file named changes that contains only these lines that were changed:

```
Through caverns measureless by man
Down by a sunless sea.
```

If the no-copy option is in effect (using the -n option on the sed command line), then the same effect can be achieved with the command

```
sed -n -e 's/to/by/p' poem > changes
```

If your command file script contains the line

```
s/[\.,;?:]/*P&*/gp
```

then the command

```
sed -n -f script poem
```

produces the output

```
A stately pleasure dome decree*P:*
Where Alph*P,* the sacred river*P,* ran
Down to a sunless sea*P.*
```

If the g flag is not used, the substitution takes effect only on the first instance of the pattern in a given line. For example, the command

```
sed -n -e '/X/s/an/AN/p' poem
```

causes the substitution to occur only on the first instance of an

```
In XANadu did Kubla Khan
```

Input/output command summary

Table 6-3 provides a summary of input and output commands for the sed stream editor.

Table 6-3 `sed` input and output commands

Command	Description
[*line1* [, *line2*]] p	(print.) The p command writes addressed lines to the standard output file. They are written at the time the p command is encountered, regardless of what succeeding editing commands may do to the lines.
[*line1* [, *line2*]] w *file*	(write to file.) The w command writes addressed lines to the file named by file. Exactly one space must separate the w and file. If the file previously existed, it is overwritten; if not, it is created.
	The lines are written exactly as they exist when the write command is encountered for each line, regardless of what subsequent editing commands may do to them. A maximum of ten different files may be mentioned in write commands and w flags after s commands combined.
[*line-spec*] r *file*	(read from file.) The r command reads the contents of file and appends them after the line matched by the address. Exactly one space must separate the r and file. The file is read and appended regardless of what subsequent editing commands may do to the line that matched its address.
	If r and a commands are run on the same line, the text from a commands and r commands is written to the output in the order that the commands are run. If a file mentioned by an r command cannot be opened, it is considered a null file, not an error, and no diagnostic is given.

◆ **Note** Since there is a limit to the number of files that can be opened simultaneously, care should be taken that no more than ten files be mentioned in w commands or flags. That number is reduced by one if any r commands are present (only one read file may be opened at a time). ◆

If the file `note1` has the following contents,

```
Note: Kubla Khan (more properly Kublai Khan;
1216-1294) was the grandson and most eminent
successor of Genghiz (Chingiz) Khan and founder
of the Mongol dynasty in China.
```

then the command

```
sed -e '/Kubla/r note1' poem
```

produces

```
In Xanadu did Kubla Khan
Note: Kubla Khan (more properly Kublai Khan;
1216-1294) was the grandson and most eminent
successor of Genghiz (Chingiz) Khan and founder
of the Mongol dynasty in China.
A stately pleasure dome decree:
Where Alph, the sacred river, ran
Through caverns measureless to man
Down to a sunless sea.
```

Multiple input line commands

The following three commands, all in uppercase letters, deal with pattern spaces containing embedded newline characters. They are intended principally to provide pattern matches across lines in the input. The `P` and `D` commands are equivalent to their lowercase counterparts if there are no embedded newline characters in the pattern space.

[*line1*[, *line2*]]N   Append the next input line to the current line in the pattern space. The two input lines are separated by an embedded newline character. Pattern matches may extend across embedded newline characters.

[*line1* [, *line2*]] D	Delete first part of the pattern space. Delete up to, and including, the first newline character in the current pattern space. If the pattern space becomes empty (the only newline character was the terminal newline character), read another line from the input. In any case, begin the list of editing commands again from the beginning.
[*line1* [, *line2*]] P	Print the first part of the pattern space. Print up to, and including, the first newline character in the pattern space.

Input commands

[*line1* [, *line2*]] h

Hold pattern space. The h command copies the contents of the pattern space into a hold area, destroying the previous contents of the hold area.

[*line1* [, *line2*]] H

Hold pattern space. The H command appends contents of the pattern space to contents of the hold area. Former and new contents are separated by a newline character.

[*line1* [, *line2*]] g

Get contents of hold area. The g command copies contents of the hold area into the pattern space, destroying previous contents.

[*line1* [, *line2*]] G

Get contents of hold area. The G command appends contents of the hold area to contents of the pattern space.

[*line1* [, *line2*]] x

Exchange. The x command interchanges contents of the pattern space and the hold area.

For example, if your `sed` command file contains the commands

```
1h
1s/ did.*//
1x
G
s/\n/ :/
```

when applied to the file `poem`, this produces

```
In Xanadu did Kubla Khan :In Xanadu
A stately pleasure dome decree: :In Xanadu
Where Alph, the sacred river, ran :In Xanadu
Through caverns measureless to man :In Xanadu
Down to a sunless sea. :In Xanadu
```

Control-flow commands

These commands do no editing on the input lines but control the application of commands to the lines selected by the address part.

[line1 [*, line2*]] !

(don't.) The exclamation point (!) command causes the next command (written on the same line) to be applied to those input lines not selected by the address part.

[line1 [*, line2*]] {

(grouping.) The { command causes the next set of commands to be applied (or not applied) as a block to the input lines selected by the addresses of the grouping command. The first of the commands under control of the grouping may appear on the same line as the { or on the next line. The group of commands is ended by a matching } standing on a line by itself. Groups can be nested.

: *label*

(place label.) The colon (:) command marks a place in the list of editing commands that may be referred to by b and t commands. The *label* argument may be any sequence

of eight or fewer characters. If two different colon commands have identical labels, a compile time diagnostic will be generated and no execution attempted.

[*line1* [*, line2*]] b *label*

(branch to label.) The b command causes the sequence of editing commands being applied to the current input line to be restarted immediately after the place where a colon command with the same *label* was encountered. The space between the b command and the *label* is optional. If no colon command with the same label can be found after all editing commands have been compiled, a compile time diagnostic is produced and no execution is attempted. A b command with no *label* is a branch to the end of the list of editing commands. Whatever should be done with the current input line is done, and another input line is read. The list of editing commands is restarted from the beginning on the new line.

[*line1* [*, line2*]] t*label*

(test substitutions.) The t command tests whether any successful substitutions have been made on the current input line; if so, it branches to *label*; if not, it does nothing. The flag which indicates that a successful substitution has been run is reset by reading a new input line or by executing a t command.

Additional commands

The following sed commands are important for working with single lines:

[*line-spec*] =

The equal sign (=) command writes the line number of the line matched by its address to the standard output.

[*line-spec*] q

The q command causes the current line to be written to the output (if it should be), any appended or read text to be written, and operation to be ended.

Index

O

o command, in vi 3-9, 3-17
open command, in ex 4-49
opening a file
  in ex 4-4
    multiple files 4-19 to 4-20
  in TextEditor 2-7
  in vi 3-3 to 3-5, 3-8
    multiple files 3-26
open mode, in ex 4-5, 4-32
open option, in ex 4-32
operators, combining with motion
    commands in vi 3-20 to 3-21
optimize option, in ex 4-32
options. *See also* flag options
  in ex 4-26 to 4-38
    listing 4-27
  for sed commands 6-4 (table)

P

p command
  in ed 5-7 to 5-8, 5-12, 5-27
  in sed 6-17, 6-19
  in vi 3-24
p option, in sed 6-15
paragraph, defined in vi 3-14
paragraph macro 4-32
paragraphs option, in ex 4-32
parentheses (( )), in vi 3-14
Paste command, in TextEditor 2-5
patterns, in ex, defining 4-16
pattern spaces, in sed 6-7
    commands for 6-18 to 6-19
percent sign (%)
  in ex 4-7, 4-18
  in resource fork of a TextEditor
    filename 2-2
period (.)
  in ed 5-4, 5-8 to 5-9, 5-10, 5-13,
    5-21, 5-28
  in ex 4-5, 4-7, 4-16, 4-44
  in sed 6-8
  in vi 3-25

plus sign (+)
  in ex 4-7
  in vi options 3-6
pound sign (#), in ex 4-18, 4-55
preserve command, in ex 4-39, 4-49
print command, in ex 4-6, 4-9 to
    4-10, 4-49
Print dialog box 2-18 (figure)
printer, selecting 2-18
printing
  entire document, in TextEditor 2-18
  in ex, current line 4-28
  partial document, in TextEditor 2-19
prompt, in ed 5-3
prompt option, in ex 4-32
put command, in ex 4-15, 4-49

Q

q command
  in ed 5-27
  in sed 6-21
question mark (?) 3-10
  in ed 5-28
quit command, in ex 4-42, 4-49 to
    4-50
quit ! command, in ex 4-39
quitting
  ed 5-6
  ex 4-41 to 4-42
  TextEditor 2-19
  vi 3-32
quotation marks 3-22, 3-24

R

r command
  in ed 5-10, 5-27
  in sed 6-17
  in vi 3-19, 3-19
r option 3-35
  in ex 4-32, 4-43 to 4-44
  in vi commands 3-6, 3-6
r suffix, in ex 4-13, 4-14, 4-51

read command, in ex 4-21 to 4-22,
    4-50
read ! command, in ex 4-24
readonly option 4-32
  in vi commands 3-6
read-only viewing
  in ex 4-32
  in vi 3-7 to 3-8
  with view 3-2
read permission, in TextEditor 2-7
recover command, in ex 4-39, 4-50
recovering files
  in ex 4-43 to 4-44
  in vi 3-6, 3-25, 3-35
recovering text, in ex 4-39
redraw option, in ex 4-33
regular expressions
  in ed 5-16
  in ex 4-16 to 4-17
    case settings 4-30
    searching 4-8, 4-36
    substitutions 4-12 to 4-13
  in sed 6-8 to 6-10, 6-14 to 6-15
  in vi 3-6, 3-25
relative motion, in ex 4-8
remap option, in ex 4-33
repeating commands 3-21 to 3-22 (table)
Replace command, in TextEditor 2-10 to
    2-11
replace commands, in vi 3-19 (table)
Replace Same command, in
    TextEditor 2-11
Replace Text dialog box 2-10 (figure)
replacing text. *See* editing text; search-
    and-replace; substitutions
report option, in ex 4-33
resource fork, of TexEditor files 2-2, 2-12
  saving 2-19
RETURN key
  in ex 4-6
  in vi 3-10
Revert to Saved command, in
    TextEditor 2-6
rewind command, in ex 4-20, 4-50

W

w command
  in ed 5-5, 5-28
  in sed 6-17
  in vi 3-32

X

x command
  in sed 6-19
  in ed 5-28
-x option, in vi commands 3-6
xit command, in ex 4-42, 4-44, 4-54

Y

y command, in vi 3-24, 3-24
yank command, in ex 4-15, 4-54
yank and put commands, in vi 3-23 to
      3-24 (table)

Z

z command, in ex 4-10 to 4-11, 4-54
ZZ command, in vi 3-32, 3-35

The Apple Publishing System

A/UX Text-Editing Tools was written, edited, and composed on a desktop publishing system using Apple Macintosh computers, an AppleTalk network system, Microsoft Word, and QuarkXPress. Line art was created with Adobe Illustrator. Proof pages were printed on Apple LaserWriter printers. Final pages were output directly to 70-mm film on an Electrocomp 2000 Electron Beam Recorder. PostScript®, the LaserWriter page-description language, was developed by Adobe Systems Incorporated.

Text type and display type are Apple's corporate font, a condensed version of ITC Garamond®. Bullets are ITC Zapf Dingbats®. Some elements, such as program listings, are set in Apple Courier, a fixed-width font.